

Evaluating ensemble post processing for wind power forecasts

Kaleb Phipps, Sebastian Lerch, Maria Andersson, Ralf Mikut, Veit Hagenmeyer, Nicole Ludwig

Angaben zur Veröffentlichung / Publication details:

Phipps, Kaleb, Sebastian Lerch, Maria Andersson, Ralf Mikut, Veit Hagenmeyer, and Nicole Ludwig. 2022. "Evaluating ensemble post-processing for wind power forecasts." *Wind Energy* 25 (8): 1379–1405. <https://doi.org/10.1002/we.2736>.

Nutzungsbedingungen / Terms of use:

CC BY 4.0



RESEARCH ARTICLE

WILEY

Evaluating ensemble post-processing for wind power forecasts

Kaleb Phipps¹  | Sebastian Lerch² | Maria Andersson³ | Ralf Mikut¹ |
Veit Hagenmeyer¹ | Nicole Ludwig^{1,4} 

¹Institute for Automation and Applied Informatics, Karlsruhe Institute of Technology, Karlsruhe, Germany

²Department of Economics, Karlsruhe Institute of Technology, Karlsruhe, Germany

³Department of Management and Engineering, Division of Energy Systems, Linköping University, Linköping, Sweden

⁴Cluster of Excellence "Machine Learning: New Perspectives for Science", University of Tübingen, Tübingen, Germany

Correspondence

Kaleb Phipps, Karlsruhe Institute of Technology, Hermann-von-Helmholtz-Platz 1, 76344 Eggenstein-Leopoldshafen, Germany.
Email: kaleb.phipps@kit.edu

Funding information

Deutsche Forschungsgemeinschaft, Grant/Award Numbers: EXC 2064/1, RTG 2153, SFB/TRR 165; Helmholtz-Gemeinschaft, Grant/Award Numbers: "Energy System Design", Helmholtz AI

Abstract

Capturing the uncertainty in probabilistic wind power forecasts is challenging, especially when uncertain input variables, such as the weather, play a role. Since ensemble weather predictions aim to capture the uncertainty in the weather system, they can be used to propagate this uncertainty through to subsequent wind power forecasting models. However, as weather ensemble systems are known to be biased and under-dispersed, meteorologists post-process the ensembles. This post-processing can successfully correct the biases in the weather variables but has not been evaluated thoroughly in the context of subsequent forecasts, such as wind power generation forecasts. The present paper evaluates multiple strategies for applying ensemble post-processing to probabilistic wind power forecasts. We use Ensemble Model Output Statistics (EMOS) as the post-processing method and evaluate four possible strategies: only using the raw ensembles without post-processing, a one-step strategy where only the weather ensembles are post-processed, a one-step strategy where we only post-process the power ensembles and a two-step strategy where we post-process both the weather and power ensembles. Results show that post-processing the final wind power ensemble improves forecast performance regarding both calibration and sharpness whilst only post-processing the weather ensembles does not necessarily lead to increased forecast performance.

KEYWORDS

ensemble post-processing, energy time series, probabilistic wind power forecasting

1 | INTRODUCTION

Forecasting, especially probabilistic forecasting, is essential to allow decision makers in power systems to optimally operate and maintain the grid.^{1,2} With the push towards energy systems with high shares of renewable energy sources, forecasting renewable generation, such as wind power, becomes increasingly important. Forecasting wind power is challenging, as the generation depends on weather variables, such as wind speed and temperature, in a non-linear and bounded fashion.³ Additionally, these weather variables are difficult to forecast and are usually described by Numerical Weather Prediction (NWP) models, which model the physical relationships of the atmosphere to a certain extent.

When forecasting wind power with the help of weather predictions, two models thus play a role: the NWP model whose output aims to describe the uncertainty in the weather variables and the wind power forecasting model whose output describes the uncertainty in the wind power given the weather variables. Within the first model, ensemble predictions, which are different runs of the NWP model, capture the

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2022 The Authors. *Wind Energy* published by John Wiley & Sons Ltd.

inherent uncertainty in the weather variables. These ensemble predictions are known to be biased and underdispersed; thus, post-processing the output of these models is done frequently.^{4–11} Post-processing is used to alleviate systematic biases in the model and calibrate the forecasts to past observations, and there is a large body of work discussing optimal post-processing methods for weather variables (see, e.g., the review by Vannitsem et al.¹²). Within the second model, i.e., the wind power forecast model, the uncertain input and the model uncertainty (i.e., epistemic uncertainty) contribute to the final probabilistic forecast. Whilst an optimal forecasting model could theoretically handle all systematic biases, post-processing might be useful when this is not the case.

In the wind power ensemble prediction setting, there are thus three ways in which systematic model biases can be present making post-processing useful: in the first stage concerning the weather output from NWP models, in the second stage concerning the wind power forecasting model or in both stages. Pinson and Messner³ explain the concept behind post-processing for wind power applications, both before the weather is used as input to the forecasting model and afterwards. However, they do not evaluate or compare these different approaches. Although the literature on wind power forecasting is vast (see, e.g., the review by Zhang et al.¹³), only a few examples dealing with post-processing exist, for example, Taylor et al.,¹⁴ Pinson and Henrik,¹⁵ Messner et al.¹⁶ and Worsnop et al.¹⁷ However, each of these papers only uses post-processing for the weather ensemble data, thus before the weather is used as input to the wind power forecasting model. To the best of our knowledge, no paper exists which evaluates whether post-processing at other stages benefits the probabilistic wind power forecast more.

In the present paper, we compare different post-processing strategies to determine at which stage in the wind power forecasting process the post-processing is most beneficial. We analyse whether one post-processing step can account for all previous biases by only post-processing the wind power ensembles and compare this one-step strategy to post-processing only the weather ensembles, i.e., assuming the biases from the wind power model are negligible. Lastly, we also compare one strategy where we post-process both weather and power ensembles. We use publicly available synthetic benchmark data and wind power data from two bidding zones in Sweden to evaluate these post-processing strategies. The remainder of the present paper is structured as follows. Section 2 introduces the theoretical background on the NWP model. We then describe the post-processing strategies in detail in Section 3, before evaluating these strategies in Section 4. Section 5 discusses our approach before Section 6 concludes the paper.

2 | NUMERICAL WEATHER PREDICTION MODELS

Atmospheric behaviour is chaotic and considered an unstable system which has finite, state dependent limits of predictability.^{18,19} Numerical Weather Prediction (NWP) models describe and forecast this atmospheric behaviour, and with it the weather, through solving a system of non-linear differential equations starting with the current observed atmospheric conditions. As this current state of the atmosphere cannot be fully observed at any given point in time and space, there remains some uncertainty with regard to the initial conditions of the NWP models. However, forecasts of non-linear numerical models are highly sensitive to the given initial conditions, and initial errors grow during the forecast.^{19,20} Accounting for the uncertainty in the initial conditions is therefore crucial and nowadays quantified with the help of Ensemble Prediction Systems (EPSs). EPSs generate ensemble forecast by running the NWP model several times with slightly different initial conditions, e.g., adding perturbations to the initial state. Hence, today's weather forecasts provide an inherent probabilistic uncertainty estimate in the form of ensembles of NWP. Figure 1 exemplarily shows the ensemble forecast trajectories over time for two consecutive forecast origins and a forecast horizon of one day. At each forecast origin, the EPS generates ensemble predictions for the specified forecast horizon. The forecast horizon describes the number of future time steps into which the weather is predicted. With increasing forecast horizon, the trajectories diverge, resulting in an




FIGURE 1 An illustration of two ensemble forecast trajectories over a 24 h forecast horizon. The red dotted lines indicate the forecast origins

increased uncertainty associated with the chaotic behaviour of the non-linear weather system. For more information on NWP, see, e.g., Bauer et al.¹⁹

3 | ENSEMBLE POST-PROCESSING STRATEGIES

The focus of the present paper is on evaluating and comparing different post-processing strategies for probabilistic wind power forecasts. Post-processing is performed to calibrate the forecasts to past observations such that systematic model biases can be alleviated. As the weather ensembles from the EPS are already known to be biased and underdispersed,⁴ different approaches exist in the meteorological literature to calibrate ensembles.¹² Two of the most common methods are Ensemble Model Output Statistics (EMOS) developed by Gneiting et al.⁸ and Bayesian Model Averaging (BMA) introduced by Raftery et al.²¹ Both of these approaches have been successful in post-processing various weather ensembles (see, for example, Javanshiri et al.²² and Han et al.²³), and the difference in performance between the two models has been shown to be negligible.^{23–25} However, EMOS is computationally far simpler than BMA,^{25,26} and currently, operational implementations of post-processing at weather services focus almost exclusively on EMOS.^{7,12} For these reasons, the present paper focuses on EMOS[†] when comparing post-processing strategies. In the following, we describe the EMOS post-processing method in detail, before discussing how it is used in our different post-processing strategies.

3.1 | Ensemble model output statistics

The EMOS method for ensemble post-processing, developed by Gneiting et al.,⁸ is based on non-homogenous regressions. The standard EMOS approach is designed for ensemble members that are individually distinguishable; however, the present paper uses ensemble members from the European Centre for Medium-Range Weather Forecasts (ECMWF). These ECMWF ensemble members are classified as *singular vector synoptic ensembles* and therefore exchangeable.^{27,28} Exchangeable ensembles represent equally likely future scenarios and have no distinguishing features or ordering. Thus, they are ensembles with a joint distribution function that is invariant under permutation of the arguments.²⁹ This exchangeability implies that, for example, the ensemble labelled as first ensemble member x_1 at forecast origin t is not related to the ensemble member with the same label at forecast origin $t + 1$. Given exchangeable ensembles, EMOS expresses a univariate weather quantity Y in terms of a multiple linear regression on the M ensemble members, with equal weights for each exchangeable ensemble member,⁷ i.e.,

$$Y = a + b \cdot \sum_{i=1}^M x_i + \epsilon. \quad (1)$$

Hereby, $x_1, \dots, x_M$ are exchangeable ensemble forecasts, a, b are regression coefficients and ϵ is an error term with a mean of zero. Given this deterministic forecast, we can create a probability density function or probabilistic forecast by assuming a distribution. The parameters of the distribution can then be modelled given the mean and variance of the individual ensemble members. For example, assuming a normal distribution for the corresponding weather variable, we use the regression on y as an approximation for the mean μ and approximate the variance σ^2 as a linear function of the ensemble spread

$$\sigma^2 = c + d \cdot S^2, \quad (2)$$

where S^2 is the ensemble variance and c and d are non-negative coefficients. The resulting parametric model is then given by

$$Y | x_1, \dots, x_M \sim \mathcal{N} \left(a + b \cdot \sum_{i=1}^M x_i, c + d \cdot S^2 \right). \quad (3)$$

Different variables require a different choice of distribution; for example, wind speeds are restricted to positive values and exhibit a skewed distribution, and EMOS can be easily extended to these other distributions.⁸

In order to estimate the EMOS coefficients, Gneiting et al.⁸ use a minimum Continuous Ranked Probability Score (CRPS) estimation which is based on the minimum contrast estimation approach. The CRPS is a measure for calibration and sharpness of a predictive cumulative distribution function F and is given by

[†]We also implement BMA and compare its performance to EMOS when calibrating individual weather ensembles. Through this experiment, we confirm the higher computational cost for BMA and observe that the calibration performance for individual weather ensembles was inferior to EMOS. Based on these results, we only consider the better performing and widely used EMOS method when comparing different post-processing strategies. We discuss this decision further in Appendix A.1.

$$\text{CRPS}(F, y) = \int_{\mathcal{R}} (F(y) - \mathbb{1}\{y \leq y\})^2 dy, \quad (4)$$

with y the verifying observation and $\mathbb{1}$ denoting an indicator function. Gneiting et al.⁸ show that the CRPS can be expressed as an analytical function and the EMOS coefficients that minimise the CRPS can be found through the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm.

The general aim of post-processing and using EMOS is to obtain calibrated forecasts. To check whether a forecast is calibrated, we look at the *probability integral transform* (PIT).³⁰ If F denotes a fixed, non-random predictive cumulative distribution function (CDF) for an observation Y , the PIT is the random variable $Z_F = F(Y)$. It is known that if F is continuous and $Y \sim F$, then Z_F is standard uniform. Thus, ideally, the PIT from the given forecast is uniform. In the discrete case, where we do not have a CDF but instead multiple ensemble members, the PIT can be described by the verification rank histogram.³¹ The verification rank histogram contains multiple bins formed from two ordered neighbouring ensemble members. Since in an ideal ensemble system the verifying ensembles are equally likely to fall within any of these bins, a rank histogram is also ideally uniformly distributed. Whilst based on slightly different principles, we interpret both PIT histograms and verification rank histograms in the same way.³¹ Figure 2 sketches the key information present in these histograms. If the post-processing is successful and the forecasts are calibrated, then we observe a uniformly distributed histogram. If the forecasts are underdispersed (i.e., they underestimate the true spread), then there are more observations in the outlying bins and less in the middle. When there is more mass in the middle of the histogram however, this is a sign of overdispersion (i.e., the forecast overestimates the true spread).

3.2 | Strategies

The focus of the present paper is the evaluation and comparison of different post-processing strategies for probabilistic wind power forecasts. Each strategy we evaluate starts with the raw weather ensembles, includes a forecasting model and has some form of probabilistic wind power forecast as output, which is then evaluated using the CRPS (see Equation 4). The strategies differ in their use of the above introduced EMOS post-processing. Figure 3 provides an overview of the forecasting process given the four different strategies, and we introduce each strategy in detail in the following.

The first strategy we introduce is called *Raw*. This strategy serves as our benchmark and does not include any form of post-processing. All available M ensemble members from the EPS are fed through the same forecasting model individually, resulting in an ensemble of wind power forecasts with M members. As the result of this strategy is not an analytical form of a distribution function, we cannot use the above CRPS formula to assess the quality of the forecast. Instead, we use the sample-based CRPS which relies on the empirical cumulative distribution function,

$$\hat{F}_M(y) = \frac{1}{M} \sum_{i=1}^M \mathbb{1}\{X_i \leq y\}, \quad (5)$$

and is calculated using

$$\text{CRPS}(\hat{F}_M, y) = \frac{1}{M} \sum_{i=1}^M |X_i - y| - \frac{1}{2M^2} \sum_{i=1}^M \sum_{j=1}^M |X_i - X_j|. \quad (6)$$




FIGURE 2 A sketch illustrating how to interpret PIT/ verification rank histograms. The ensembles are underdispersed if there is more mass in the outer bins (left diagram) and overdispersed if there is more mass in the central bins (right diagram). Well-calibrated ensembles are uniformly distributed (middle diagram)




FIGURE 3 An overview of the post-processing possibilities available. We identify four strategies: no post-processing (*Raw*), post-processing only the power ensembles (*One-Step-P*), post-processing only the weather ensembles (*One-Step-W*) and post-processing both weather and power ensembles (*Two-Step-WP*)

In the second strategy, *One-Step-P*, post-processing is applied once for the final wind power forecasting model. This strategy relies on the assumption that post-processing the wind power ensembles also accounts for the biases in the weather ensembles. We again start with the M raw ensemble members and feed them individually through the forecasting model. In contrast to the *Raw* strategy, the resulting ensemble of wind power forecasts is post-processed. Thus, we fit the ensemble of wind power forecasts to the historical wind power generation using a rolling EMOS approach, where the parameters are estimated every day on the basis of the past 40 days. Following Gneiting and Katzfuss,³² we assume a truncated normal distribution for the wind power using $\mathcal{N}_{[0,\infty)}^+(\mu, \sigma^2)$, with location $\mu = a + b\bar{x}$ and scale $\sigma^2 = c + dS^2$ as an affine function of the ensemble variance $S^2 = \frac{1}{M} \sum_{i=1}^M (x_i - \bar{x})^2$. The resulting wind power forecast can then be evaluated using the CRPS for truncated normal distributions

$$\text{CRPS}(F_{\text{TN}}, y) = \sigma \left[\frac{y - \mu}{\sigma} \Phi(\mu/\sigma) (2\Phi((x - \mu)/\sigma) + \Phi(\mu/\sigma) - 2) + 2\varphi((y - \mu)/\sigma) \Phi(\mu/\sigma) - \frac{1}{\sqrt{\pi}} \Phi\left(\sqrt{2}\mu/\sigma\right) \right] [\Phi(\mu/\sigma)]^{-2}. \quad (7)$$

As an alternative, we also consider a gamma distribution for the wind power. For simplicity, the main paper focuses on the results using the truncated normal distribution, with the almost identical gamma distribution results presented in the Section A.2.

In the third strategy, *One-Step-W*, we post-process only the weather ensembles. Thus, instead of using the raw ensembles, as for the two previous strategies, we post-process the ensembles and then draw M independent samples from each of the resulting calibrated weather distributions. Whilst independent sampling does ignore possible dependency structures that exist between weather variables, Phipps et al.³³ show that restoring these structures with methods such as Ensemble Copula Coupling (ECC) has no noticeable effect on the results. Therefore, we only consider a simple independent sampling method for this paper. Furthermore, we select M samples to replicate the number of raw ensembles available. Overall, the *One-Step-W* strategy relies on the assumption that all biases in the wind power forecast can be eliminated by accounting for the biases in the weather ensembles. For the post-processing, each weather variable is considered separately, and for each weather variable, we use EMOS with a rolling training window to estimate the parameters for the appropriate distributions. We use the distributions as suggested by Gneiting,⁷ thus a normal distribution for temperature, a normal distribution for the u - and v -components of wind and a truncated normal distribution or a gamma distribution for wind speed.[‡] Given these distributions, we draw M independent random samples from each distribution to form weather calibrated ensemble members. We then feed these sampled ensemble members through the forecasting model, resulting in an ensemble of wind power forecasts whose accuracy can be assessed with the sample-based CRPS from Equation 6.

In the fourth and last strategy, *Two-Step-WP*, both the weather ensembles and the ensemble of wind power forecasts are post-processed. This strategy relies on the assumption that neither of the one-step approaches can sufficiently account for all biases in the models and data and the forecasts should be post-processed at all stages in the forecasting process. Thus, the two one-step strategies are coupled together, where we

[‡]Again, the gamma and truncated normal distribution delivered similar performance. For simplicity, we focus on the results from the truncated normal distribution in the main paper and report full results in Section A.2.

first apply the procedure described for *One-Step-W* to the weather ensembles and then *One-Step-P* to the ensemble of wind power forecasts. The distribution choices, parameters and CRPS calculations remain the same as for the individual strategies.

3.3 | Forecasting model

We aim to evaluate ensemble post-processing at different stages in the forecasting process; thus, we also need an appropriate wind power forecasting model. This forecasting model should represent the relationship between the different weather variables and the wind power and, in contrast to a manufacturers' wind power curve model, can be estimated regardless of the turbine type and for aggregated wind power values (e.g., the sum of wind power from a large region with an unknown number of wind turbines). Since we focus on the comparison of different post-processing strategies and not raw forecast accuracy, we want the models to be simple, common in the literature and lightweight with respect to parameter optimisation and computation time. Research has shown that post-processing generally improves regression results, even when deep learning methods are applied (see, e.g., Kuleschov et al.³⁴), and therefore, we expect the post-processing performance to be similar independent of the model chosen. We therefore choose one linear model, namely, a linear regression,³⁵ as a simple benchmark model and one more complex non-linear model, namely, a random forest.³⁶ Furthermore, to show the robustness of our results, we additionally implement an alternative linear regression model based on Zhang and Wang³⁷ and a simple artificial neural network as an alternative non-linear forecasting model. In the following, we describe the linear regression and random forest models as well as how we measure the forecasting accuracy. The additional linear regression and neural network models are presented in Section A.2.

The linear regression model which we use to forecast the wind power given the weather input can be described with

$$y_{t+h} = \beta_0 + \alpha y_{t+h-24} + \sum_{k=1}^K \beta_k W_{t+h}^k + \sum_{j=1}^J \gamma_j D_{t+h}^j + \varepsilon_{t+h}, \quad (8)$$

where y_t is the dependent variable, which in this case is the wind power; y_{t-24} is the actual wind power a day before; W^k are weather time series, such as wind speed, the u-component of wind, the v-component of wind, surface pressure and temperature; and D^j are dummy variables, such as the season, the month and the year. The models are fitted for each forecast horizon $h = h_1, \dots, h_H$ with $h_H \leq 24$, using actual historical weather data in order to describe the real relationship among the variables and remove any bias fitting on historical weather forecasts or ensembles could introduce. Each ensemble $x_1 \dots x_M$ from the EPS is then used in a separate prediction run for each forecast horizon to generate an ensemble of wind power predictions with the previously fitted regression coefficients

$$\hat{y}_{t+h}(x_1, \dots, x_M) = \hat{\beta}_0 + \hat{\alpha} y_{t+h-24} + \sum_{k=1}^K \hat{\beta}_k \hat{W}_{t+h}^k(x_1, \dots, x_M) + \sum_{j=1}^J \hat{\gamma}_j \hat{D}_{t+h}^j + \hat{\varepsilon}_t. \quad (9)$$

In order to better forecast non-linear dependencies, we also implement a random forest. Random forests are a statistical learning method and are additionally to their ability to model non-linearity also robust to errors when unnecessary predictors are included. Random forests create several de-correlated regression trees to form a collection of solutions. These de-correlated trees are fitted by selecting a subset of features at each candidate split. The final prediction is then the average over all regression trees. We use a random forest with 500 trees and the same input variables as for the linear regression to predict the wind power. Giving the random forest the full set of input variables available, specifically multiple weather features, allows it to automatically identify the most important features and hereby generate accurate forecasts. The parameters of the random forest are, equivalently to the linear regression, estimated using the historical values, whilst the probabilistic ensemble forecasts are generated by running the forest on each ensemble member.

The forecasting strategy is the same, regardless of the specific forecasting model. Given the forecast origin, we use weather forecasts $\hat{W}_{t+h}^k$ from the EPS to predict the wind power $\hat{y}_{t+h}$ for each of the forecasts horizons h (between 3 and 24 h and determined by the temporal resolution of the EPS). Additionally to the weather variables, we also include historical wind power generation from 24 h before the prediction time and dummy features $\hat{D}_{t+h}^j$, such as the time of day or month. Due to the fact that the historical weather observations are implicitly included in the NWP model and the calibration process relies on these observations, we do not include historical weather information as a specific input for our prediction models. All forecasting models are implemented in R, and for the random forests, we use the *randomForest* package.[§] We apply the forecast strategy explained above for every post-processing strategy shown in Section 3.2, such that the forecasting strategy does not influence our comparison. To assess the performance, we calculate the average CRPS over all time steps $t = 1, \dots, N$ in the test set

[§]<https://cran.r-project.org/web/packages/randomForest/randomForest.pdf>

$$\text{CRPS} = \frac{1}{N} \sum_{t=1}^N \text{CRPS}(F_t, Y_t). \quad (10)$$

As we are mainly interested in a fair comparison of our post-processing strategies, we also calculate the CRPS Skill Score (CRPSS). This score measures the improvement in CRPS when compared to a given benchmark model. In our case, the CRPSS is calculated for all post-processing strategies with respect to the *Raw* strategy as the benchmark. We report the CRPSS as a percentage, where positive results indicate an improvement.

4 | EVALUATION

Given the four ensemble post-processing strategies introduced in the previous section, we want to evaluate their performance by comparing the accuracy, sharpness and calibration of the resulting wind power forecasts. Before we report the results, we introduce the two data sets we use: an open source benchmark data set (Section 4.1) and a data set with real wind power generation data from two bidding zones in Sweden (Section 4.2). In Section 4.3, we present the results of different post-processing strategies using only the truncated normal distribution, whilst the results of further experiments with different distributions are reported in Section A.2.

4.1 | Benchmark data

Due to the lack of open source wind power data for specific wind parks, we generate benchmark wind power using the *Renewables.ninja* API.[†] Staffell and Pfenninger³⁸ verify that the simulation and bias-corrections implemented in the *Renewables.ninja* API are capable of reproducing accurate wind power time series. We replicate two real German wind power parks, one onshore and one offshore. Table 1 shows the parameters which we use, where we selected turbines with similar characteristics than those installed using the wind turbine database.[#]

We use The International Grand Global Ensemble (TIGGE) archive[‡] to access open source ensemble weather data. TIGGE archive is a result of *The Observing System Research and Predictability Experiment* which aimed to combine ensemble forecasts from leading forecast centres to improve probabilistic forecasting capabilities.³⁹ The archive includes a limited sample of the ECMWF ensembles from October 2006 until the present. The limitations are placed on the available forecast horizons (only in steps of 6 h instead of 3 h in the licenced version of the EPS), the number of weather variables (e.g., wind speed and wind components are only available at a height of 10 m and not also on 100 m) and a reduced spatial resolution compared to the operational EPS. TIGGE archive is however publicly accessible and data can be downloaded through the MARS API. We use weather data for the same locations as the synthetically generated wind power data (see Table 1). We include the parameters temperature at 2 m above ground, surface pressure, u-Component and v-Component of wind at 10 m and wind speed, for the period from February 2017 until August 2018. The limited time span of data available is due to damaged tapes in TIGGE archive which affected all ensemble data, and the aforementioned period is the longest available with continuous weather data at the time of writing. For the ground truth historical weather data, we use the ERA5 reanalysis data.⁴⁰ This data are accessed via the Copernicus Climate Data Store (CDS) API.^{**} Here, the same locations and identical parameters are included. When working with the benchmark dataset, we use the entire year 2017 for training and the remainder of the data (01.2018–08.2018) for the evaluation.

In order to allow replication of the results and further development of the methods presented in the current paper, the ensemble weather data and the synthetic wind power time series are made available through GitHub (<https://github.com/KIT-IAI/EvaluatingEnsemblePostProcessing>). Due

TABLE 1 Configuration parameters used to generate the wind power time series with the *Renewable.ninja* API

	Onshore benchmark	Offshore benchmark
Coordinates	51.0° N, 10.5° E	54.5° N, 6.0° E
Time span	01-02-2017 to 31-08-2018	01-02-2017 to 31-08-2018
Capacity	130 MW	400 MW
Turbine height	95 m	90 m
Turbine type	Vestas V90 2000	Gamesa G128 5000
Similar real windpark	Windfeld Wangenheim-Hochheim-Ballstädt-Westhausen	BARD Offshore I

[†]www.renewables.ninja.

[#]<https://en.wind-turbine-models.com/turbines>

[‡]<https://apps.ecmwf.int/datasets/data/tigge/>

^{**}<https://cds.climate.copernicus.eu/home>

to licencing constraints, the ERA5 reanalysis data must be downloaded separately. We provide instructions for downloading and formatting the reanalysis data in the repositories documentation.

4.2 | Swedish data set

The second data set we use contains data from the Swedish electricity system. The Swedish electricity system is divided into four sub-areas or bidding zones.⁴¹ The division of the electricity system into sub-areas has several purposes.⁴² One purpose is to create regional price differences between the sub-areas. This is an incentive for cost-effective further development of the electricity system, as new power plants will be built where there is a shortage of electricity. Another purpose is that state-owned Svenska Kraftnät, which operates the national grid, should receive indications about where the national grid needs to be strengthened to be able to transfer enough electricity to the other sub-areas. A third purpose is to comply with EU legislation and thereby facilitate the continued integration of the Swedish electricity system with the European electricity market.

In the present paper, we focus on the area contained in bidding zones 3 and 4, for two reasons. First, the large majority of wind power generation occurs in these two bidding zones (see Figure 4). Second, these are the two bidding zones in Sweden that are sometimes faced with a lack of electricity.⁴² Whilst in northern Sweden the supply of electricity is normally greater than the demand, transmission capacity between north and south of Sweden is not always sufficient to transfer the demanded electricity, and this can lead to bottlenecks.⁴²

Weather data for bidding zones 3 and 4 consist of the ECMWF EPS (Molteni et al.²⁷) and also the ERA5 reanalysis data C3S.⁴⁰ We use the EPS as the foundation for probabilistic forecasting methods and again use the ERA5 reanalysis data as the ground truth for the post-processing. Table 2 summarises the key aspects of the data. The weather data available are in a grid-based format. This means that atmospheric models are used to create a NWP for certain geographical grid-points on earth. Since we are not considering a single wind park but looking at the aggregated wind power generation for each bidding zone, we considered all of these grid points in the form of a weighted average. The weighted average is calculated as follows; first, each of the data points is sorted into the appropriate bidding zone based on the geographical specifications




FIGURE 4 A map of Sweden with the four bidding zones shown through the border lines. The distribution of wind turbines is indicated by the blue circles. The figure is adapted from Olauson and Bergkvist⁴³

TABLE 2 A summary of the key characteristics of the weather data available for the use case in Sweden

Characteristic	Notes
Temporal dimension	05.01.2015 to 31.08.2019
Spatial dimension	11° E–19.5° E 54° N–62.5° N
Spatial resolution	Grid resolution of 0.25° × 0.25°
Forecast time	Forecasts for up to 24 h ahead made at 00:00:00
Forecast horizon	One step ahead forecasts for 3, 6, 9, 12, 15, 21 and 24 h
Weather variables	100 m u-component of wind, 100 m v-component of wind, 100 m wind speed, 2 m temperature, surface pressure

(see Figure 4); then, a weighted average of every point is calculated. Due to a lack of accurate location data for various wind turbines in Sweden, a rudimentary weighted average method is used; areas with a high concentration of wind turbines are given double weighting, whilst those areas with a lower concentration only receive a standard weight. As seen in Figure 4, the doubly weighted areas include a central area in bidding zone 3 and the coastal areas in bidding zone 4.

The wind power generation data are available through the open source transparency platform which is operated by the European Network of Transmission System Operators (ENTSO-E).⁴⁴ The transparency platform provides aggregated onshore wind power generation data at an hourly resolution from the 05.01.2009 until present. These data are aggregated for each bidding zone in Sweden (zones 1–4), where we consider bidding zones 3 and 4 in the present paper. We use data from 2015–2017 for the training of our models and then from 01.2018–08.2019 for the evaluation. Table 3 provides a statistical summary of the wind power generation data collected. The data collected are the raw wind power generation in Megawatt and therefore affected by structural changes such as an increase in capacity, outages due to maintenance and upgrades to wind turbines.

4.3 | Results

In this section, we present the evaluation results given the different post-processing strategies and the forecasting models introduced in Section 3. We start with the results for the benchmark data set (Section 4.3.1) before reporting the results for the Swedish data set (Section 4.3.2).

TABLE 3 A statistical summary of the wind power generation data for both benchmark data sets and for bidding zones 3 and 4 in Sweden

	Onshore benchmark	Offshore benchmark	Bidding zone 3	Bidding zone 4
Minimum	0.13 MW	1.2 MW	3 MW	2 MW
1st quartile	11.31 MW	125.6 MW	278 MW	167 MW
Median	23.53 MW	236.8 MW	568 MW	352 MW
Mean	29.12 MW	223.9 MW	670.96 MW	447.55 MW
3rd quartile	40.04 MW	333.2 MW	967 MW	664 MW
Maximum	128.83 MW	385.6 MW	2246 MW	1463 MW

TABLE 4 Summary of mean CRPS on the test data for the benchmark data sets and for all forecast horizons

Data set		6 h	12 h	18 h	24 h
Onshore Benchmark	Linear Raw	3.91	5.60	5.67	6.08
	Linear <i>One-Step-P</i>	3.53	4.88	5.71	5.76
	Linear <i>One-Step-W</i>	3.81	5.26	5.51	5.90
	Linear <i>Two-Step-WP</i>	3.48	4.97	5.68	5.77
	Random Forest Raw	3.98	5.34	5.68	5.84
	Random Forest <i>One-Step-P</i>	3.61	4.63	5.84	5.38
	Random Forest <i>One-Step-W</i>	3.97	4.53	5.78	5.83
	Random Forest <i>Two-Step-WP</i>	3.67	4.34	5.84	5.70
Offshore Benchmark	Linear Raw	18.98	23.51	25.46	24.11
	Linear <i>One-Step-P</i>	17.92	22.23	24.85	23.44
	Linear <i>One-Step-W</i>	18.74	23.31	24.51	23.88
	Linear <i>Two-Step-WP</i>	17.30	22.08	24.28	23.49
	Random Forest Raw	20.77	23.80	25.00	25.54
	Random Forest <i>One-Step-P</i>	19.47	22.15	24.24	23.83
	Random Forest <i>One-Step-W</i>	21.45	23.22	25.41	24.93
	Random Forest <i>Two-Step-WP</i>	20.47	22.43	25.56	24.08

Note: The best prediction for each strategy, forecast horizon and model is highlighted in bold.

4.3.1 | Benchmark data

We first perform the different post-processing strategies on the benchmark data set and evaluate the calibration and sharpness based on the above introduced CRPS and PIT/ verification rank histogram for both forecasting models with varying forecast horizon. Table 4 summarises the mean CRPS for the onshore and offshore benchmark data set with the best value highlighted in bold. We can see that for almost every model and every forecast horizon either the *One-Step-P* or *Two-Step-WP* post-processing strategy perform the best. Figure 5 plots the CRPS for each




FIGURE 5 The CRPS skill score plotted against the forecast horizon on the test data for the onshore benchmark (top figure) and the offshore benchmark (bottom figure). Positive values indicate an improvement over the *Raw* strategy in per cent. In general, post-processing leads to an improvement in forecasting performance, although only post-processing the weather has little impact or performs slightly worse




FIGURE 6 An overview of the post-processing results for a forecast horizon of 12 h on the onshore benchmark data set. Although each individual weather variable is calibrated before the wind power forecast, this does not result in a calibrated wind power ensemble. In order to achieve a calibrated wind power ensemble, post-processing after the wind power forecast is required

post-processing strategy across all forecast horizons. We see that post-processing almost always leads to an improvement in CRPS when compared to the *Raw* strategy. The improvement depends on the forecast horizon and the data set, but is typically between 5% and 15%.

We next take a closer look at the calibration, i.e., PIT/ verification rank histograms. For brevity's sake, we focus in the following on the onshore data set and pick a forecast horizon of 12 h (note that the other data set and forecast horizons behave similarly and are therefore not presented in detail). The effect of the various post-processing strategies on the calibration is shown in Figure 6, where we plot the different verification rank histograms and PITs at each forecasting step. The raw weather ensembles for wind speed, temperature, and the u-component of wind

TABLE 5 Summary of mean CRPS on the test data for the use case in Sweden for all forecast horizons

Data Set		3h	6h	9h	12h	15h	18h	21h	24h
Bidding Zone 3	Linear Raw	57.84	78.46	81.03	79.75	86.20	89.05	89.06	93.33
	Linear <i>One-Step-P</i>	47.30	55.99	62.01	67.39	67.31	64.36	67.17	68.00
	Linear <i>One-Step-W</i>	64.51	86.61	90.05	87.27	96.89	99.55	97.95	98.99
	Linear <i>Two-Step-WP</i>	47.08	55.32	62.04	65.56	66.86	63.70	67.05	65.83
	Random Forest Raw	60.71	71.31	69.72	67.74	73.39	82.01	83.67	88.09
	Random Forest <i>One-Step-P</i>	46.08	50.28	55.09	55.96	55.17	53.44	57.83	58.86
	Random Forest <i>One-Step-W</i>	69.96	79.95	78.21	76.50	82.77	90.38	91.47	91.15
	Random Forest <i>Two-Step-WP</i>	45.20	49.92	55.48	54.27	55.36	53.95	57.34	57.24
Bidding Zone 4	Linear Raw	41.75	49.91	50.93	63.17	63.85	53.27	51.46	51.21
	Linear <i>One-Step-P</i>	36.17	41.60	43.15	49.68	53.51	47.86	45.92	43.69
	Linear <i>One-Step-W</i>	41.91	51.20	51.85	57.17	61.52	52.86	51.24	48.86
	Linear <i>Two-Step-WP</i>	35.78	41.24	43.38	50.61	54.43	47.50	47.15	43.14
	Random Forest Raw	36.97	41.33	40.74	50.67	51.60	45.54	42.26	39.41
	Random Forest <i>One-Step-P</i>	34.33	38.25	39.30	44.60	48.02	44.28	41.72	38.18
	Random Forest <i>One-Step-W</i>	36.81	42.91	41.94	46.51	51.16	46.24	43.16	38.31
	Random Forest <i>Two-Step-WP</i>	33.52	38.25	39.21	43.95	48.08	44.74	42.04	37.33

Note: The best prediction for each strategy and each forecast model is highlighted in bold.




FIGURE 7 The CRPS skill score plotted against the forecast horizon on the test data for both bidding zone 3 (top figure) and bidding zone 4 (bottom figure) in Sweden. Positive values indicate an improvement over the *Raw* strategy in per cent. Post-processing the final power ensemble, either directly or as part of a two-step process, clearly improves the forecast, whilst only post-processing the weather ensemble leads to worse or similar forecast performance

measured at 10 m show a clear positive bias and are slightly underdispersed. The v-component of wind measured at 10 m does not show the same positive bias but is also underdispersed. After the weather ensembles are post-processed, the resulting PIT histograms appear almost uniform. Interestingly, the v-component of wind now appears to be slightly overdispersed which indicates an overcorrection in the post-processing. The wind power ensemble forecasts generated with the raw weather ensembles are underdispersed, but there is almost no improvement when the post-processed weather ensembles are used. The ensemble forecast from One-Step-P appears uniform, whilst the PIT histogram for Two-Step-WP also indicates a slight overcorrection with overdispersion present.

4.3.2 | Swedish data

This section focuses on the evaluation of the Swedish data set. We again evaluate the mean CRPS (a summary of these values is shown in Table 5) for each model and each forecast horizon. We also consider the CRPSS in Figure 7, where the improvement in CRPS compared to the *Raw* strategy is plotted for each bidding zone in Sweden against the forecast horizon. Considering bidding zone 3, we see that processing the final power ensemble (*One-Step-P* or *Two-Step-WP*) leads to vast improvements in CRPS, with a CRPSS ranging between 20% and 40%. However, the *One-Step-W* strategy, which only processes the weather ensembles, leads to a decrease in performance of approximately 10%. For bidding zone 4, the results are similar. Calibrating the power ensemble again results in clear CRPS improvements, whilst only calibrating the weather ensembles leads to slightly worse performance, or performance very similar to the *Raw* strategy (excluding a forecast horizon of 12 h). Furthermore, we observe that although the mean CRPS for different forecasting models (linear regression and random forests) are noticeably different, the CRPSS results are very similar. This can be observed on both the bidding zone 3 and bidding zone 4 data set.

Additionally to the mean CRPS, we also consider box plots of the CRPS to help assess the effect of post-processing on the spread of the forecast, i.e., the sharpness. Figure 8 shows these box plots of all CRPS for a forecast horizon of 12 h, exemplary for bidding zone 3 in Sweden. A smaller box indicates a sharper forecast. Figure 8a details the performance of the linear model, and it can be seen that both *One-Step-P* and *Two-Step-WP* improve the spread of the CRPS quite noticeably. *One-Step-W* has little effect on the median CRPS but does lead to a larger spread, suggesting that this model struggles with values that are difficult to predict.

For the random forests (Figure 8b), the CRPS is smaller (note the scale of the y axis). Apart from an overall improvement in CRPS, the sharpness results are similar to the linear models. There is no clear improvement in median, but both *One-Step-P* and *Two-Step-WP* strategies show a noticeably smaller spread in CRPS values. This smaller spread is an indication that the post-processing results in a better approximation of those values that are normally difficult to predict (i.e., result in a large CRPS value), even if the median accuracy is not noticeably different.

5 | DISCUSSION

The results above show that post-processing the ensembles leads to more accurate wind power forecasts. As the two strategies which include post-processing the final wind power ensemble (*One-Step-P* and *Two-Step-WP*) outperform the others (*One-Step-W* and *Raw*), this final post-processing seems crucial. Only post-processing the weather (*One-Step-W*) does not improve the forecast performance regarding both calibration and sharpness; in some cases, it even leads to a decrease in performance. A possible explanation for this behaviour is that post-processing only




FIGURE 8 A boxplot of all CRPS values for a forecast horizon of 12 h for bidding zone 3 in Sweden. In (A), the forecast using a linear model is shown and in (B) the forecast using a random forest. Although the post-processing does not show large improvement in the median, it does reduce the spread

the weather neglects other bias sources for the wind power model. We train the models on historical data such that they learn the true relationship among the variables; however, this makes it harder for the forecasting model to properly propagate the uncertainty from the set of ensembles through the model, as they are not designed to do that. Thus, despite the weather ensembles being well-calibrated after post-processing, this information seems to not be accurately used by the forecasting model.

It is also worth noting that to enable comparisons across all strategies, we only consider a sample size of 51, i.e., the number of raw ensembles, when drawing from the calibrated weather ensembles. Although it is possible to vary the number of independently drawn samples, there are two reasons why we conclude this will not have a noticeable effect on the results. First, the considered sample size of 51 is large enough to accurately approximate the estimated distribution, and, second, increasing the number of samples will also increase computational complexity and could lead to poorer performance in subsequent calibration steps. However, a systematic analysis of the effect of sample size could be interesting in future work.

When we compare *One-Step-P* and *Two-Step-WP*, both strategies perform well, and in most cases, there is no noticeable difference between them. The key factor that differentiates the two strategies is the number of post-processing steps required. When we apply the *Two-Step-WP* strategy, we need to post-process all weather inputs (in the case of the current paper this is four, but in more complicated models, this number may increase). Once this post-processing is complete, we generate a wind power forecast and then are again required to process the resulting ensembles. The *One-Step-P* strategy involves only one post-processing step, irrespective of how many weather inputs the model includes. Although EMOS' computational complexity is low and multiple post-processing steps are still feasible, *One-Step-P* is a more computationally efficient strategy that achieves similar forecasting accuracy. Overall, the post-processing of wind power forecasts does increase forecasting accuracy, independent of the forecasting model, with the last post-processing step (in this case the wind power) having the biggest effect.

6 | CONCLUSION

The present paper evaluates at which stage ensemble post-processing is the most useful for wind power forecasts. Although Pinson and Messner³ introduced ensemble post-processing for wind power applications, they did not evaluate which post-processing strategy is most effective. We, therefore, identify four post-processing strategies that can be applied: (1) no post-processing (*Raw*), (2) a one-step strategy with post-processing of the resulting wind power ensemble (*One-Step-P*), (3) a one-step strategy with post-processing of the weather ensembles (*One-Step-W*) and (4) a two-step strategy with both post-processing of the weather ensembles and the resulting wind power ensembles (*Two-Step-WP*). These strategies are evaluated on two data sets with a linear regression and a random forest used as forecasting models. Results show that post-processing generally improves performance, specifically when the wind power ensemble is post-processed. *One-Step-P* and *Two-Step-WP* deliver similar results in terms of CRPS and CRPSS performance, but since it requires significantly fewer post-processing steps, the *One-Step-P* strategy is preferred.

Whilst initial results show that post-processing the final wind power ensemble leads to improvements in forecasting accuracy, there are numerous aspects that should be investigated further. First, EMOS is limited in that it assumes a parametric distribution in the calibration process and wind power generation and other weather variables do not always perfectly follow such a distribution. Therefore, non-parametric post-processing methods that do not rely on the assumption of a parametric probability distribution should also be considered. Furthermore, the current paper focuses on a one-step-ahead forecast for either a single location or a broad region. This removes the need to account for spatial and temporal dependencies. Ensemble Copula Coupling (ECC) could be used to account for these dependencies, and future work should focus on multi-step ahead and multi-location forecasts. It could also be interesting to systematically analyse the effect of sample size on forecasting performance in future work. Finally, the current methods use traditional forecasting methods which do not appear to be able to learn the uncertainty in the data. Therefore, approaches that learn this uncertainty in the weather ensembles and directly propagate it through to a wind power forecast should be investigated.

ACKNOWLEDGEMENTS

This work is funded by the German Research Foundation (DFG) as part of the Research Training Group 2153 'Energy Status Data – Informatics Methods for its Collection, Analysis and Exploitation' and by the Helmholtz Association's Initiative and Networking Fund through Helmholtz AI, the Helmholtz Association under the Program 'Energy System Design' and the Joint Initiative 'Energy System Design - A Contribution of the Research Field Energy'. Nicole Ludwig acknowledges financial support by the DFG under Germany's Excellence Strategy – EXC number 2064/1 – Project number 390727645. Sebastian Lerch acknowledges support by the DFG through SFB/TRR 165 'Waves to Weather' and by the Vector Stiftung through the Young Investigator Group 'Artificial Intelligence for Probabilistic Weather Forecasting'. The research detailed in the current paper was based on data from the ECMWF obtained through an academic licence for research purposes. The authors would like to thank Jon Olafson for his help creating Figure 4 and Marian Turowski for the useful discussions and his continued personal support throughout the development of this paper.

PEER REVIEW

The peer review history for this article is available at <https://publons.com/publon/10.1002/we.2736>.

DATA AVAILABILITY STATEMENT

The ensemble weather data for the onshore and offshore benchmarks that support the findings of this study are openly available via GitHub (<https://github.com/KIT-IAI/EvaluatingEnsemblePostProcessing>) and were adapted from The International Grand Global Ensemble (TIGGE) archive (<https://apps.ecmwf.int/datasets/data/tigge/>). The wind power generation data for the onshore and offshore benchmarks are generated with the Renewables.ninja API (<https://www.renewables.ninja/>) and also openly available via GitHub (<https://github.com/KIT-IAI/EvaluatingEnsemblePostProcessing>). The ensemble weather data for Sweden that further support the findings of this study are generated from the Ensemble Prediction System (EPS) from European Centre for Medium-Range Weather Forecasts (ECMWF). These data were downloaded with a license from ECMWF and are provided with the Creative Commons Attribution 4.0 International (CC BY 4.0) license. We adapt the data from the ECMWF archive and make it openly available thanks to ECMWF via GitHub (<https://github.com/KIT-IAI/EvaluatingEnsemblePostProcessing>) in line with the CC BY 4.0 licence. The wind power generation data for Sweden are openly available via the open source transparency platform, operated by the European Network of Transmission System Operators (ENTSO-E).⁴⁴ The observational weather data used in this study are ERA5 reanalysis data and openly available via the Copernicus Climate Data Store (CDS) (<https://cds.climate.copernicus.eu/home>).

ORCID

Kaleb Phipps  <https://orcid.org/0000-0002-9197-1739>

Nicole Ludwig  <https://orcid.org/0000-0003-3230-8918>

REFERENCES

- Ordiano JAG, Waczowicz S, Hagenmeyer V, Mikut R. Energy forecasting tools and services. *WIREs Data Mining Knowl Discov*. 2017;8(2):e1235.
- Appino RR, Ordiano JAG, Mikut R, Faulwasser T, Hagenmeyer V. On the use of probabilistic forecasts in scheduling of renewable energy sources coupled to storages. *Appl Energy*. 2018;210:1207-1218.
- Pinson P, Messner JW. Chapter 9 - application of postprocessing for renewable energy. In: Vannitsem S, Wilks DS, Messner JW, eds. *Statistical Post-processing of Ensemble Forecasts*; Elsevier; 2018:241-266.
- Thorarinsdottir TL, Gneiting T. Probabilistic forecasts of wind speed: ensemble model output statistics by using heteroscedastic censored regression. *J R Stat Soc: Ser A (Stat Soc)*. 2010;173(2):371-388.
- Schefzik R, Thorarinsdottir TL, Gneiting T. Uncertainty quantification in complex simulation models using ensemble copula coupling. *Stat Sci*. 2013;28(4):616-640.
- Schefzik R. Ensemble calibration with preserved correlations: unifying and comparing ensemble copula coupling and member-by-member post-processing. *Quart J R Meteorol Soc*. 2017;143(703):999-1008.
- Gneiting T. Calibration of medium-range weather forecasts, Technical Memorandum 719. ECMWF; 2014.
- Gneiting T, Raftery AE, Westveld AH, Goldman T. Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation. *Month Weather Rev*. 2005;133(5):1098-1118.
- Gilbert C, Messner JW, Pinson P, Trombe P-J, Verzijlbergh R, Dorp P, Jonker H. Statistical post-processing of turbulence-resolving weather forecasts for offshore wind power forecasting. *Wind Energy*. 2020;23(4):884-897.
- Sweeney C, Lynch P. Adaptive post-processing of short-term wind forecasts for energy applications. *Wind Energy*. 2011;14(3):317-325.
- Bossavy A, Girard R, Kariniotakis G. Forecasting ramps of wind power production with numerical weather prediction ensembles. *Wind Energy*. 2013;16(1):51-63.
- Vannitsem S, Bremnes JB, Demaeyer J, et al. Statistical postprocessing for weather forecasts: review, challenges, and avenues in a big data world. *Bull Am Meteorol Soc*. 2021;102(3):E681-E699.
- Zhang Y, Wang J, Wang X. Review on probabilistic forecasting of wind power generation. *Renew Sustain Energy Rev*. 2014;32:255-270.
- Taylor JW, McSharry PE, Buizza R. Wind power density forecasting using ensemble predictions and time series models. *IEEE Trans Energy Convers*. 2009;24(3):775-782.
- Pinson P, Madsen H. Ensemble-based probabilistic forecasting at horns rev. *Wind Energy*. 2009;12(2):137-155.
- Messner JW, Zeileis A, Broecker J, Mayr GJ. Probabilistic wind power forecasts with an inverse power curve transformation and censored regression. *Wind Energy*. 2014;17(11):1753-1766.
- Worsnop RP, Scheuerer M, Hamill TM, Lundquist JK. Generating wind power scenarios for probabilistic ramp event prediction using multivariate statistical post-processing. *Wind Energy Sci*. 2018;3(1):371-393.
- Lorenz EN. Deterministic nonperiodic flow. *J Atmos Sci*. 1963;20(2):130-141.
- Bauer P, Thorpe A, Brunet G. The quiet revolution of numerical weather prediction. *Nature*. 2015;525(7567):47-55.
- Poincaré H, Maitland F. *Science and Method*: Courier Corporation; 2003.
- Raftery AE, Gneiting T, Balabdaoui F, Polakowski M. Using bayesian model averaging to calibrate forecast ensembles. *Month Weather Rev*. 2005;133(5):1155-1174.
- Javanshiri Z, Fathi M, Mohammadi A. Comparison of the BMA and EMOs statistical methods for probabilistic quantitative precipitation forecasting. *Meteorol Appl*. 2021;28:e1974.
- Han K, Choi J, Kim C. Comparison of statistical post-processing methods for probabilistic wind speed forecasting. *Asia-Pacific J Atmos Sci*. 2018;54(1):91-101.

24. Baran S, Hornyí A, Nemoda D. Comparison of BMA and EMOs statistical calibration methods for temperature and wind speed ensemble weather prediction. *Időjárás*. 2013;118(3):217-241. 2014.
25. Schulz B, Lerch S. Machine learning methods for postprocessing ensemble forecasts of wind gusts: a systematic comparison. *Month Weather Rev*. 2021;150:235-257.
26. Vannitsem S, Wilks DS, Messner J. *Statistical Postprocessing of Ensemble Forecasts*: Elsevier; 2018.
27. Molteni F, Buizza R, Palmer TN, Petroliagis T. The ECMWF ensemble prediction system: methodology and validation. *Quart J R Meteorol Soc*. 1996;122(529):73-119.
28. Fraley C, Raftery AE, Gneiting T. Calibrating multimodel forecast ensembles with exchangeable and missing members using bayesian model averaging. *Month Weather Rev*. 2010;138(1):190-202.
29. Brcker J, Kantz H. The concept of exchangeability in ensemble forecasting. *Nonlinear Processes Geophys*. 2011;18(1):1-5.
30. Angus JE. The probability integral transform and related results. *SIAM Rev*. 1994;36(4):652-654.
31. Hamill TM. Interpretation of rank histograms for verifying ensemble forecasts. *Month Weather Rev*. 2001;129(3):550-560.
32. Gneiting T, Katzfuss M. Probabilistic forecasting. *Ann Rev Stat Appl*. 2014;1(1):125-151.
33. Phipps K, Ludwig N, Hagenmeyer V, Mikut R. Potential of ensemble copula coupling for wind power forecasting. In: Proceedings 30. Workshop Computational Intelligence Berlin; 26-27. November 2020, Vol. 26 KIT Scientific Publishing; 2020:87.
34. Kuleshov V, Fenner N, Ermon S. Accurate uncertainties for deep learning using calibrated regression. In: Proceedings of Machine Learning Research, Vol. 80. PMLR; 2018:2796-2804. <https://proceedings.mlr.press/v80/kuleshov18a.html>
35. Wu Y-K, Hong J-S. A literature review of wind forecasting technology in the world. In: 2007 IEEE Lausanne POWERTECH, Proceedings; 2007: 504-509.
36. Lydia M, Kumar SS. A comprehensive overview on wind power forecasting. In: 2010 Conference Proceedings IPEC, IEEE; 2010.
37. Zhang Y, Wang J. K-nearest neighbors and a kernel density estimator for gefcom2014 probabilistic wind power forecasting. *Int J Forecast*. 2016;32(3): 1074-1080.
38. Staffell I, Pfenninger S. Using bias-corrected reanalysis to simulate current and future wind power output. *Energy*. 2016;114:1224-1239.
39. Swinbank R, Kyouda M, Buchanan P, et al. The TIGGE project and its achievements. *Bull Am Meteorol Soc*. 2016;97(1):49-67.
40. Copernicus Climate Change Service (C3S). Era5: Fifth generation of ecmwf atmospheric reanalyses of the global climate Edited by Copernicus Climate Change Service Climate Data Store (CDS). Copernicus Climate Change Service; 2017.
41. Energy Market Inspectorate. Elområden i sverige analys av utvecklingen och konsekvenserna på marknaden, eir 2012:06. El R2012:06; 2012.
42. Elområden.se. Din guide till nya elmarknaden. <https://elomraden.se/om>. Accessed on 2019-12-16; 2011.
43. Olason J, Bergkvist M. Modelling the swedish wind power production using MERRA reanalysis data. *Renew Energy*. 2015;76:717-725.
44. European Network of Transmission System Operators (ENTSO-E). Transparency platform (tp). Online Document. <https://transparency.entsoe.eu/>; 2019.
45. Fraley C, Raftery AE, Gneiting T, Sloughter JM. Ensemblebma: An r package for probabilistic forecasting using ensembles and bayesian model averaging, WASHINGTON UNIV SEATTLE DEPT OF STATISTICS; 2007.

How to cite this article: Phipps K, Lerch S, Andersson M, Mikut R, Hagenmeyer V, Ludwig N. Evaluating ensemble post-processing for wind power forecasts. *Wind Energy*. 2022;25(8):1379-1405. doi:[10.1002/we.2736](https://doi.org/10.1002/we.2736)

APPENDIX A

In this appendix, we first discuss the performance of Bayesian Model Averaging (BMA) when post-processing weather ensembles in Section A.1. In Section A.2, we then report further results of experiments not included in the main paper that serve to demonstrate the robustness of our results. These experiments include using an alternative linear regression and a neural network as forecasting models, considering alternative distributions for wind speed and power and showing that the selected distributions are a reasonable fit for the data.

A.1 | Bayesian model averaging as an alternative post-processing method

As discussed in Section 3, one common alternative to EMOS is BMA, introduced by Raftery et al.²¹ Although BMA is computationally more expensive than EMOS,^{25,26} we initially implement both post-processing methods and compare their performance regarding the calibration of weather ensembles. We implement BMA with the R package *ensembleBMA*,⁴⁵ selecting the same probability distributions for weather variables as chosen for EMOS.

Figure 9 exemplarily shows the calibration performance when using BMA for a forecast horizon of 12 h for the onshore benchmark data set, the same scenario used in Figure 6. Based on this figure, it is clear that the performance of BMA is inferior to EMOS. Although temperature is only slightly underdispersed, all other variables are poorly calibrated with wind speed being overdispersed and both the u- and v-components of wind very underdispersed. The calibration performance of EMOS for the same scenario (see Figure 6) is far superior. We discover similar results for the other data sets and forecast horizons. We believe this is due to the computational complexity of BMA which caused convergence




FIGURE 9 PIT Histograms showing the calibration of weather ensembles using BMA for the onshore benchmark data set with a forecast horizon of 12 h. In comparison to EMOS, the ensembles are poorly calibrated

problems on our data set. Based on these initial results, we select EMOS as the best performing post-processing method on our data set and only consider EMOS for the evaluation of our developed post-processing strategies.

A.2 | Robustness of results

To show the robust nature of our approach, we perform numerous further experiments. First, we implement both an alternative linear regression and a neural network as further forecasting methods, before also including the full results when considering different distributions for wind power and wind speed. Finally, we show that the selected probability distributions are a reasonable fit for our data and therefore a suitable choice for EMOS.

A.2.1 | Alternative linear regression model

To discuss the results of the alternative linear regression model proposed by Zhang and Wang,³⁷ we first describe the linear regression used before showing the results of the different post-processing strategies.

The linear regression proposed by Zhang and Wang³⁷ only considers wind speed and the hour of the day as regressors. The linear regression model can be described as

$$y_{t+h} = \beta_0 + \beta_1 S_{t+h} + \beta_2 (S_{t+h})^2 + \beta_3 (S_{t+h})^3 + \beta_4 \cos\left(\frac{2\pi}{24} H_{t+h}\right) + \beta_5 \sin\left(\frac{2\pi}{24} H_{t+h}\right) + \beta_6 \cos\left(\frac{4\pi}{24} H_{t+h}\right) + \beta_7 \sin\left(\frac{4\pi}{24} H_{t+h}\right) + \varepsilon_{t+h}, \quad (\text{A1})$$

where y_t is the dependent variable, which in this case is the wind power; S_{t+h} is the forecasted wind speed for the considered forecast horizon; H_{t+h} is the hour of the day for the considered forecast horizon and $\beta_i, i=0, \dots, 7$ are the regression coefficients. As observed by Zhang and Wang,³⁷ the sine and cosine terms account for diurnal periodicities through a Fourier series considering the time of the day. We fit and apply this alternative linear regression model in the same way as the linear regression model described in Section 3.3.

The CRPS skill score plotted against the forecast horizon for the different post-processing strategies is shown for the benchmark data sets in Figure 10 and for both bidding zones in Sweden in Figure 11. These results again highlight that post-processing the final wind power ensemble is the crucial step. The *One-Step-P* and *Two-Step-WP* strategies, which post-process the wind power ensemble, almost always improve forecasting performance on all data sets. The only exceptions are for a 24 h forecast horizon in the offshore benchmark and a 21 h forecast performance in bidding zone 4, which deliver forecast performance similar to the *Raw* strategy. As with the linear regression model presented in Section 3.3, post-processing the weather ensembles does not generally lead to improvements. In general, the improvements achieved through post-processing with the alternative linear model are similar to the improvements achieved with other forecasting models and therefore show the robust nature of our approach.




FIGURE 10 The CRPS skill score plotted against the forecast horizon on the test data for the alternative linear regression on the onshore benchmark (top figure) and the offshore benchmark (bottom figure). Positive values indicate an improvement over the *Raw* strategy in per cent. In general, post-processing leads to an improvement in forecasting performance




FIGURE 11 The CRPS skill score plotted against the forecast horizon on the test data for the alternative linear regression on both bidding zone 3 (top figure) and bidding zone 4 (bottom figure) in Sweden. Positive values indicate an improvement over the *Raw* strategy in per cent. Post-processing the final power ensemble, either directly or as part of a two-step process, almost always clearly improves the forecast, whilst only post-processing the weather ensemble leads to worse or similar forecast performance




FIGURE 12 The CRPS skill score plotted against the forecast horizon on the test data for the neural networks on the onshore benchmark (top figure) and the offshore benchmark (bottom figure). Positive values indicate an improvement over the Raw strategy in per cent. In general, post-processing leads to an improvement in forecasting performance

A.2.2 | Neural network as an alternative non-linear forecasting model

To discuss the results using neural networks as an additional forecasting model, we first describe the selected neural network architecture and then show the results of the different post-processing strategies.

We select a simple feed forward neural network for our evaluation and test multiple configurations before choosing a configuration with two hidden layers of 10 and 7 neurons, respectively. This network architecture is selected because it is the simplest we found that still returns accurate forecasts. Given this architecture, we train the neural network with the resilient backpropagation algorithm. The chosen activation function is a hyperbolic tangent, and the input features are the same as those selected for the linear regression and random forest models (see Section 3.3). The parameters (i.e., weights) are fitted using the actual historical weather data, and each ensemble member is passed through the network to get an ensemble of wind power forecasts.

The CRPS skill score plotted against the forecast horizon for the different post-processing strategies is shown for the benchmark data sets in Figure 12 and for both bidding zones in Sweden in Figure 13. With these results, we again see that post-processing the final wind power ensemble is the crucial step. Both the *One-Step-P* and *Two-Step-WP* strategies, which post-process the wind power ensemble, improve performance on the benchmark and the Swedish data sets. Only post-processing the weather ensembles, however, does not necessarily lead to improvements. Although *One-Step-W* leads to improved performance on the benchmark data set, this strategy performs worse on the real world Swedish data sets. Furthermore, we note that the improvements through the post-processing strategies are similar to the improvements achieved with other forecasting models.

A.2.3 | Considering different distributions for wind speed and wind power

From the literature, we identify both the truncated normal distribution and gamma distribution as possible candidates for modelling wind speed and wind power. To ensure robust, results we consider both distributions and report the CRPS values when using different combinations of the two distributions here. Following Gneiting,⁷ we implement a gamma distribution using $\text{Gamma}(\alpha, \beta) = F_{\alpha, \beta}$, with $\alpha = a + b_1x_1 + \dots + b_Mx_M$ and $\beta = c + dm_{\text{ENS}}$ with $m_{\text{ENS}} = a + b \sum_{m=1}^M x_m$. The following tables (Tables 6–9) show that there is no noticeable performance difference when using different distributions. Additionally, we also report the results of models that do not include weather forecasts as input and are based solely on historical wind power values.




FIGURE 13 The CRPS skill score plotted against the forecast horizon on the test data for the neural networks on both bidding zone 3 (top figure) and bidding zone 4 (bottom figure) in Sweden. Positive values indicate an improvement over the Raw strategy in per cent. Post-processing the final power ensemble, either directly or as part of a two-step process, clearly improves the forecast, whilst only post-processing the weather ensemble leads to worse or similar forecast performance

A.2.4 | Fitting distributions

To verify that the choice of distributions for the EMOS calibration is valid, we fit the selected distributions for both the weather variables and resulting wind power generation to the data. Figures 14–17 show the fit of the selected distributions: a truncated normal and gamma distribution for wind power and wind speed, and Gaussian distributions for all other weather variables.

It is noticeable that although the fits are not perfect, the selected distributions are a good approximation for the data. Some of the empiric distributions are difficult to represent with parametric distributions, and therefore, these results further motivate non-parametric approaches, which we plan to consider in future work.

TABLE 6 Summary of mean CRPS on the test data for the onshore benchmark for all forecast horizons

Strategy	6 h	12 h	18 h	24 h
Linear No Weather	7.24	11.38	11.42	10.97
Linear Raw	3.91	5.60	5.67	6.08
Linear <i>One-Step-P</i> (Gamma)	3.42	4.68	5.60	5.48
Linear <i>One-Step-P</i> (T-Normal)	3.53	4.88	5.71	5.76
Linear <i>One-Step-W</i> (Gamma)	3.80	5.29	5.66	5.96
Linear <i>One-Step-W</i> (T-Normal)	3.81	5.26	5.51	5.90
Linear <i>Two-Step-WP</i> (W: Gamma, P: Gamma)	3.52	4.96	5.63	5.78
Linear <i>Two-Step-WP</i> (W: Gamma, P: T-Normal)	3.42	4.96	5.61	5.72
Linear <i>Two-Step-WP</i> (W: T-Normal, P: Gamma)	3.50	4.87	5.64	5.84
Linear <i>Two-Step-WP</i> (W: T-Normal, P: T-Normal)	3.48	4.97	5.68	5.77
Alternative Linear No Weather	13.45	15.73	12.45	11.46
Alternative Linear Raw	4.59	5.79	6.21	5.93
Alternative Linear <i>One-Step-P</i> (Gamma)	3.67	4.12	5.78	5.43
Alternative Linear <i>One-Step-P</i> (T-Normal)	3.81	4.35	5.91	5.53
Alternative Linear <i>One-Step-W</i> (Gamma)	4.73	4.32	6.31	5.82
Alternative Linear <i>One-Step-W</i> (T-Normal)	4.73	4.19	6.38	5.91
Alternative Linear <i>Two-Step-WP</i> (W: Gamma, P: Gamma)	4.01	4.22	5.86	5.52
Alternative Linear <i>Two-Step-WP</i> (W: Gamma, P: T-Normal)	3.99	4.22	5.99	5.63
Alternative Linear <i>Two-Step-WP</i> (W: T-Normal, P: Gamma)	3.98	3.93	5.89	5.56
Alternative Linear <i>Two-Step-WP</i> (W: T-Normal, P: T-Normal)	3.85	4.04	6.15	5.89
Random Forest No Weather	9.52	12.18	10.41	10.31
Random Forest Raw	3.98	5.34	5.68	5.84
Random Forest <i>One-Step-P</i> (Gamma)	3.31	4.21	5.67	5.23
Random Forest <i>One-Step-P</i> (T-Normal)	3.61	4.63	5.84	5.38
Random Forest <i>One-Step-W</i> (Gamma)	3.85	4.44	5.82	5.83
Random Forest <i>One-Step-W</i> (T-Normal)	3.97	4.53	5.78	5.83
Random Forest <i>Two-Step-WP</i> (W: Gamma, P: Gamma)	3.34	4.08	5.73	5.45
Random Forest <i>Two-Step-WP</i> (W: Gamma, P: T-Normal)	3.57	4.30	5.86	5.52
Random Forest <i>Two-Step-WP</i> (W: T-Normal, P: Gamma)	3.41	4.05	5.73	5.42
Random Forest <i>Two-Step-WP</i> (W: T-Normal, P: T-Normal)	3.67	4.34	5.84	5.70
Neural Network No Weather	7.80	11.34	10.15	10.31
Neural Network Raw	7.16	10.42	10.19	10.19
Neural Network <i>One-Step-P</i> (Gamma)	6.05	8.57	8.56	7.83
Neural Network <i>One-Step-P</i> (T-Normal)	6.19	9.12	8.56	8.02
Neural Network <i>One-Step-W</i> (Gamma)	5.85	9.49	8.96	9.43
Neural Network <i>One-Step-W</i> (T-Normal)	5.90	9.60	9.38	9.44
Neural Network <i>Two-Step-WP</i> (W: Gamma, P: Gamma)	5.46	8.61	8.02	8.01
Neural Network <i>Two-Step-WP</i> (W: Gamma, P: T-Normal)	5.54	9.00	7.71	7.76
Neural Network <i>Two-Step-WP</i> (W: T-Normal, P: Gamma)	5.53	8.76	8.21	7.89
Neural Network <i>Two-Step-WP</i> (W: T-Normal, P: T-Normal)	5.44	8.99	7.71	7.65

Note: The distribution used for post-processing the wind speed (W) and wind power (P) is shown in parentheses.

TABLE 7 Summary of mean CRPS on the test data for the offshore benchmark for all forecast horizons

Strategy	6 h	12 h	18 h	24 h
Linear No Weather	35.64	54.01	66.41	68.15
Linear Raw	18.98	23.51	25.46	24.11
Linear One-Step-P (Gamma)	17.97	22.52	24.95	23.69
Linear One-Step-P (T-Normal)	17.92	22.23	24.85	23.44
Linear One-Step-W (Gamma)	18.33	23.54	24.15	23.81
Linear One-Step-W (T-Normal)	18.74	23.31	24.51	23.88
Linear Two-Step-WP (W: Gamma, P: Gamma)	17.64	22.98	23.84	24.76
Linear Two-Step-WP (W: Gamma, P: T-Normal)	17.21	22.36	24.03	23.38
Linear Two-Step-WP (W: T-Normal, P: Gamma)	18.19	22.67	24.10	25.10
Linear Two-Step-WP (W: T-Normal, P: T-Normal)	17.30	22.08	24.28	23.49
Alternative Linear No Weather	69.48	70.90	72.09	69.61
Alternative Linear Raw	21.25	24.25	26.25	24.60
Alternative Linear One-Step-P (Gamma)	19.94	23.48	25.86	24.68
Alternative Linear One-Step-P (T-Normal)	19.85	23.12	25.34	25.10
Alternative Linear One-Step-W (Gamma)	20.97	24.12	25.93	23.88
Alternative Linear One-Step-W (T-Normal)	21.24	23.93	26.49	24.65
Alternative Linear Two-Step-WP (W: Gamma, P: Gamma)	20.19	24.21	26.42	25.79
Alternative Linear Two-Step-WP (W: Gamma, P: T-Normal)	20.22	23.67	25.29	24.94
Alternative Linear Two-Step-WP (W: T-Normal, P: Gamma)	20.48	23.81	27.14	26.40
Alternative Linear Two-Step-WP (W: T-Normal, P: T-Normal)	20.96	24.13	26.09	25.92
Random Forest No Weather	51.37	58.67	64.36	64.75
Random Forest Raw	20.77	23.80	25.00	25.54
Random Forest One-Step-P (Gamma)	18.91	21.97	25.03	24.54
Random Forest One-Step-P (T-Normal)	19.47	22.15	24.24	23.83
Random Forest One-Step-W (Gamma)	21.26	23.20	25.07	24.79
Random Forest One-Step-W (T-Normal)	21.45	23.22	25.41	24.93
Random Forest Two-Step-WP (W: Gamma, P: Gamma)	20.31	21.94	25.23	25.26
Random Forest Two-Step-WP (W: Gamma, P: T-Normal)	20.24	22.27	24.49	23.64
Random Forest Two-Step-WP (W: T-Normal, P: Gamma)	20.33	22.51	25.67	25.13
Random Forest Two-Step-WP (W: T-Normal, P: T-Normal)	20.47	22.43	25.56	24.08
Neural Network No Weather	37.99	54.47	64.61	68.55
Neural Network Raw	42.61	55.57	49.78	47.01
Neural Network One-Step-P (Gamma)	38.18	46.36	46.47	42.79
Neural Network One-Step-P (T-Normal)	38.19	45.67	44.03	43.71
Neural Network One-Step-W (Gamma)	37.80	49.98	44.11	41.22
Neural Network One-Step-W (T-Normal)	38.58	51.33	45.52	42.19
Neural Network Two-Step-WP (W: Gamma, P: Gamma)	36.89	43.32	42.87	39.69
Neural Network Two-Step-WP (W: Gamma, P: T-Normal)	36.07	42.74	40.47	38.90
Neural Network Two-Step-WP (W: T-Normal, P: Gamma)	38.13	43.60	44.01	40.38
Neural Network Two-Step-WP (W: T-Normal, P: T-Normal)	38.12	43.20	42.24	39.70

Note: The distribution used for post-processing the wind speed (W) and wind power (P) is shown in parentheses.

TABLE 8 Summary of mean CRPS on the test data for the use case of bidding zone 3 in Sweden for all forecast horizons

Strategy	3 h	6 h	9 h	12 h	15 h	18 h	21 h	24 h
Linear No Weather	68.25	125.24	169.95	206.38	228.61	244.21	255.13	273.54
Linear Raw	57.84	78.46	81.03	79.75	86.20	89.05	89.06	93.33
Linear One-Step-P (Gamma)	45.87	55.41	62.07	67.34	70.53	64.36	67.97	67.45
Linear One-Step-P (T-Normal)	47.30	55.99	62.01	67.39	67.31	64.36	67.17	68.00
Linear One-Step-W (Gamma)	63.98	86.41	89.92	87.47	95.43	98.89	97.16	98.64
Linear One-Step-W (T-Normal)	64.51	86.61	90.05	87.27	96.89	99.55	97.95	98.99
Linear Two-Step-WP (W: Gamma, P: Gamma)	46.03	55.53	62.90	67.13	69.74	64.54	67.28	67.65
Linear Two-Step-WP (W: Gamma, P: T-Normal)	46.42	54.69	61.06	66.47	67.40	64.72	65.97	65.77
Linear Two-Step-WP (W: T-Normal, P: Gamma)	46.73	54.88	62.44	66.65	69.04	62.50	68.34	66.76
Linear Two-Step-WP (W: T-Normal, P: T-Normal)	47.08	55.32	62.04	65.56	66.86	63.70	67.05	65.83
Alternative Linear No Weather	276.01	282.10	280.85	287.40	292.19	290.27	279.83	280.57
Alternative Linear Raw	82.32	82.86	76.98	70.64	73.19	81.16	82.95	83.68
Alternative Linear One-Step-P (Gamma)	59.74	58.48	56.85	57.52	56.18	55.63	60.11	56.36
Alternative Linear One-Step-P (T-Normal)	59.95	59.01	57.42	58.53	57.43	55.85	61.02	57.44
Alternative Linear One-Step-W (Gamma)	97.98	93.85	86.63	80.77	84.45	90.63	91.10	89.64
Alternative Linear One-Step-W (T-Normal)	98.29	94.21	86.78	80.51	85.04	89.73	90.44	89.02
Alternative Linear Two-Step-WP (W: Gamma, P: Gamma)	57.43	56.73	56.79	55.95	56.55	55.19	60.48	58.29
Alternative Linear Two-Step-WP (W: Gamma, P: T-Normal)	58.22	58.27	56.81	56.76	56.73	56.47	60.51	57.52
Alternative Linear Two-Step-WP (W: T-Normal, P: Gamma)	57.62	57.29	55.46	55.96	55.87	55.07	60.19	56.55
Alternative Linear Two-Step-WP (W: T-Normal, P: T-Normal)	58.74	58.99	56.67	57.58	57.48	56.41	60.88	57.72
Random Forest No Weather	170.90	192.13	211.51	235.68	252.47	252.25	251.61	260.93
Random Forest Raw	60.71	71.31	69.72	67.74	73.39	82.01	83.67	88.09
Random Forest One-Step-P (Gamma)	43.54	48.30	52.78	55.29	54.46	52.92	57.36	57.11
Random Forest One-Step-P (T-Normal)	46.08	50.28	55.09	55.96	55.17	53.44	57.83	58.86
Random Forest One-Step-W (Gamma)	70.05	79.57	78.28	76.32	82.51	91.18	91.08	92.40
Random Forest One-Step-W (T-Normal)	69.96	79.95	78.21	76.50	82.77	90.38	91.47	91.15
Random Forest Two-Step-WP (W: Gamma, P: Gamma)	43.00	48.10	53.28	54.17	54.95	53.46	56.25	57.52
Random Forest Two-Step-WP (W: Gamma, P: T-Normal)	44.84	49.53	55.36	54.19	54.58	54.02	56.75	56.08
Random Forest Two-Step-WP (W: T-Normal, P: Gamma)	42.83	48.83	53.37	54.44	54.34	52.51	55.76	55.73
Random Forest Two-Step-WP (W: T-Normal, P: T-Normal)	45.20	49.92	55.48	54.27	55.36	53.95	57.34	57.24
Neural Network No Weather	60.11	111.15	150.67	185.20	202.59	210.39	224.67	233.61
Neural Network Raw	54.05	78.40	70.00	76.04	100.56	73.69	73.16	75.79
Neural Network One-Step-P (Gamma)	48.26	61.59	58.07	63.21	76.20	61.80	67.13	67.65
Neural Network One-Step-P (T-Normal)	48.38	64.54	59.89	66.47	78.67	64.66	69.14	68.11
Neural Network One-Step-W (Gamma)	56.59	78.18	72.74	71.19	98.05	72.66	72.87	73.20
Neural Network One-Step-W (T-Normal)	56.75	78.24	72.64	71.44	99.37	74.27	72.51	75.06
Neural Network Two-Step-WP (W: Gamma, P: Gamma)	48.31	61.04	59.85	63.69	76.67	63.02	69.13	67.87
Neural Network Two-Step-WP (W: Gamma, P: T-Normal)	48.56	64.67	60.48	64.27	77.11	63.77	68.66	66.32
Neural Network Two-Step-WP (W: T-Normal, P: Gamma)	48.09	62.30	60.30	63.38	75.97	63.69	68.40	67.69
Neural Network Two-Step-WP (W: T-Normal, P: T-Normal)	48.70	64.57	60.81	65.20	76.92	64.61	68.46	67.70

Note: The distribution used for post-processing the wind speed (W) and wind power (P) is shown in parentheses.

TABLE 9 Summary of mean CRPS on the test data for the use case of bidding zone 4 in Sweden for all forecast horizons

Strategy	3 h	6 h	9 h	12 h	15 h	18 h	21 h	24 h
Linear No Weather	49.39	89.66	120.37	142.60	154.74	165.48	172.85	181.03
Linear Raw	41.75	49.91	50.93	63.17	63.85	53.27	51.46	51.21
Linear <i>One-Step-P</i> (Gamma)	35.70	42.39	45.62	54.83	58.10	50.11	49.84	47.72
Linear <i>One-Step-P</i> (T-Normal)	36.17	41.60	43.15	49.68	53.51	47.86	45.92	43.69
Linear <i>One-Step-W</i> (Gamma)	41.78	51.18	52.17	57.57	61.95	52.73	51.68	49.12
Linear <i>One-Step-W</i> (T-Normal)	41.91	51.20	51.85	57.17	61.52	52.86	51.24	48.86
Linear <i>Two-Step-WP</i> (W: Gamma, P: Gamma)	35.44	42.47	45.94	54.84	58.20	50.14	49.22	47.74
Linear <i>Two-Step-WP</i> (W: Gamma, P: T-Normal)	36.08	40.73	42.56	49.52	53.96	48.05	45.62	43.49
Linear <i>Two-Step-WP</i> (W: T-Normal, P: Gamma)	35.31	42.24	45.74	54.04	57.33	49.50	49.64	46.95
Linear <i>Two-Step-WP</i> (W: T-Normal, P: T-Normal)	35.78	41.24	43.38	50.61	54.43	47.50	47.15	43.14
Alternative Linear No Weather	184.17	189.06	196.37	202.03	197.75	190.11	186.76	187.90
Alternative Linear Raw	42.92	44.07	42.82	54.38	55.16	44.77	40.33	38.30
Alternative Linear <i>One-Step-P</i> (Gamma)	39.76	40.06	40.17	48.42	50.34	43.93	41.02	37.65
Alternative Linear <i>One-Step-P</i> (T-Normal)	39.68	40.71	40.93	48.11	51.96	44.59	41.61	38.13
Alternative Linear <i>One-Step-W</i> (Gamma)	42.95	46.43	44.15	51.19	55.29	45.58	40.60	38.03
Alternative Linear <i>One-Step-W</i> (T-Normal)	43.37	46.75	44.18	50.98	55.58	45.98	41.08	38.17
Alternative Linear <i>Two-Step-WP</i> (W: Gamma, P: Gamma)	38.16	40.60	39.50	48.42	50.72	43.77	41.47	38.36
Alternative Linear <i>Two-Step-WP</i> (W: Gamma, P: T-Normal)	38.99	41.41	40.37	47.80	50.26	44.40	41.41	37.82
Alternative Linear <i>Two-Step-WP</i> (W: T-Normal, P: Gamma)	38.91	40.44	39.50	48.32	50.54	43.47	41.41	37.58
Alternative Linear <i>Two-Step-WP</i> (W: T-Normal, P: T-Normal)	39.15	40.42	40.11	48.34	51.39	44.39	41.73	37.90
Random Forest No Weather	116.32	129.71	148.64	166.77	171.40	168.17	169.95	175.11
Random Forest Raw	36.97	41.33	40.74	50.67	51.60	45.54	42.26	39.41
Random Forest <i>One-Step-P</i> (Gamma)	33.34	37.24	38.29	44.51	47.28	43.09	40.58	37.88
Random Forest <i>One-Step-P</i> (T-Normal)	34.33	38.25	39.30	44.60	48.02	44.28	41.72	38.18
Random Forest <i>One-Step-W</i> (Gamma)	36.63	42.63	41.89	46.57	51.59	45.69	43.07	38.57
Random Forest <i>One-Step-W</i> (T-Normal)	36.81	42.91	41.94	46.51	51.16	46.24	43.16	38.31
Random Forest <i>Two-Step-WP</i> (W: Gamma, P: Gamma)	32.47	37.03	37.82	43.25	48.01	43.12	41.11	37.42
Random Forest <i>Two-Step-WP</i> (W: Gamma, P: T-Normal)	33.71	37.91	38.76	43.33	47.90	44.73	41.72	37.32
Random Forest <i>Two-Step-WP</i> (W: T-Normal, P: Gamma)	32.44	37.31	37.87	43.74	46.81	43.13	40.73	37.25
Random Forest <i>Two-Step-WP</i> (W: T-Normal, P: T-Normal)	33.52	38.25	39.21	43.95	48.08	44.74	42.04	37.33
Neural Network No Weather	49.75	90.51	121.89	150.22	157.56	163.37	168.79	174.76
Neural Network Raw	47.52	48.68	49.41	58.64	58.55	61.21	55.49	51.71
Neural Network <i>One-Step-P</i> (Gamma)	38.92	44.02	45.03	57.58	54.49	49.75	47.89	44.22
Neural Network <i>One-Step-P</i> (T-Normal)	40.26	44.24	45.80	56.82	56.79	51.21	47.55	45.02
Neural Network <i>One-Step-W</i> (Gamma)	70.80	92.27	105.05	100.64	98.54	80.15	88.04	87.43
Neural Network <i>One-Step-W</i> (T-Normal)	51.76	68.62	77.67	78.28	74.47	61.77	65.74	65.53
Neural Network <i>Two-Step-WP</i> (W: Gamma, P: Gamma)	43.55	53.51	51.32	61.07	58.26	55.46	49.56	50.08
Neural Network <i>Two-Step-WP</i> (W: Gamma, P: T-Normal)	45.79	53.14	52.60	59.16	57.91	55.56	49.88	52.13
Neural Network <i>Two-Step-WP</i> (W: T-Normal, P: Gamma)	40.89	51.27	50.19	56.03	55.45	52.22	47.70	48.15
Neural Network <i>Two-Step-WP</i> (W: T-Normal, P: T-Normal)	43.45	50.78	49.57	56.19	54.98	52.43	48.17	49.43

Note: The distribution used for post-processing the wind speed (W) and wind power (P) is shown in parentheses.




FIGURE 14 Fitted distributions for wind power and all weather variables for the offshore benchmark data set




FIGURE 15 Fitted distributions for wind power and all weather variables for the onshore benchmark data set




FIGURE 16 Fitted distributions for wind power and all weather variables for bidding zone 3 in Sweden




FIGURE 17 Fitted distributions for wind power and all weather variables for bidding zone 4 in Sweden