

Article

Probabilistic Forecasting of Wind and Solar Farm Output

John Boland ^{*,†}  and Sleiman Farah [†]

Industrial AI Research Centre, UniSA STEM, University of South Australia, Adelaide, SA 5001, Australia; sleiman.farah@unisa.edu.au

* Correspondence: john.boland@unisa.edu.au

† These authors contributed equally to this work.

Abstract: Accurately forecasting the output of grid connected wind and solar systems is critical to increasing the overall penetration of renewables on the electrical network. This includes not only forecasting the expected level, but also putting error bounds on the forecast. The National Electricity Market (NEM) in Australia operates on a five minute basis. We used statistical forecasting tools to generate forecasts with prediction intervals, trialing them on one wind and one solar farm. In classical time series forecasting, construction of prediction intervals is rudimentary if the error variance is constant—termed homoscedastic. However, if the variance changes—either conditionally as with wind farms, or systematically because of diurnal effects as with solar farms—the task is much more complicated. The tools were trained on segments of historical data and then tested on data not used in the training. Results from the testing set showed good performance using metrics, including Coverage and Interval Score. The methods used can be adapted to various time scales for short term forecasting.

Keywords: solar farms; wind farms; probabilistic forecasting; prediction interval; homoscedastic; autoregressive moving average (ARMA) models; exponential smoothing; heteroscedastic; autoregressive conditional heteroscedastic (ARCH) effect



Citation: Boland, J.; Farah, S.

Probabilistic Forecasting of Wind and Solar Farm Output. *Energies* **2021**, *14*, 5154. <https://doi.org/10.3390/en14165154>

Academic Editors: Tapas Mallick and Abu-Siada Ahmed

Received: 23 June 2021

Accepted: 18 August 2021

Published: 20 August 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

In 2018, the Australian Renewable Energy Agency (ARENA) announced a funding scheme. The focus was on improving '5-min ahead' forecasts for wind and solar farms operating in the National Electricity Market (NEM).

The way the NEM works is that there are three types of generators. Scheduled generators submit a bid stack every five minutes detailing how much electricity they can supply in the subsequent five minutes at each of ten price bands from AUD-1000–AUD14000 per MW. They are termed price makers. Renewable energy generators with capacity between 30 and 100 MW are termed semi-scheduled generators. They do not submit bids, but can be curtailed and are termed price takers. Generators of any type under 30 MW are non-scheduled and cannot be curtailed.

In partnership with the Australian Energy Market Operator (AEMO), ARENA is seeking to demonstrate that wind and solar farms can provide more accurate forecasts of their output into AEMO's central dispatch system. At present, wind and solar farms can be disadvantaged if their available output does not match AEMO's forecast. If the forecasts are too low, wind and solar farms are restricted in how much electricity they can be paid to produce. If forecasts are too high, the wind or solar farm may be obliged to pay for the cost of stabilising the frequency, which increases the price of electricity and this is ultimately passed on to consumers.

The Solar Power Ensemble Forecaster (SPEF) project was one of the ones funded by ARENA to assist the solar farms in developing their forecast models. In the SPEF project, researchers from University of South Australia (UniSA), University of New South Wales (UNSW), The Commonwealth Scientific and Industrial Research Organisation (CSIRO), and Industrial Monitoring and Control (IMC) utilised an ensemble of forecast models,

which are combined using machine learning to generate a single forecast output. Each of the models utilises a different technology or technique, such as using sky-camera and satellite images to analyse cloud evolution and motion to forecast electrical power output.

The SPEF project [1] produced substantial results including:

- Significant increases in forecast performance were achieved, as compared to AEMO's incumbent forecasting system;
- Significant causer pays fees savings were achieved from operational forecasts submitted to AEMO at all four sites analysed. In the existing causer pays approach, if participants deviate from their dispatch trajectory in a way that makes the frequency worse (i.e., they are below target when the frequency is low, or vice versa), then a penalty is determined for the participant.

However, one expects that the results, especially for the value proposition, could be enhanced if probabilistic forecasting of wind and solar farms was employed. This would be especially true in the present climate in Australia, where, increasingly, the operators of these facilities have onsite storage systems. Large-scale renewable energy plants with storage capabilities for power of 5 MW or more correspond to scheduled generators, which sell their electricity through the spot market. As a consequence, the managers of large-scale farms integrated with an energy storage must schedule the output profile of their systems and submit bids into the market. The short term bid stack is submitted every five minutes for the following five minute interval. One can imagine that if one were to have a probabilistic forecast available, one could examine the width of the prediction interval and make strategic decisions about how to act.

Therefore, a solar or wind farm with 5 MW storage could decide if the prediction interval is narrow to set their bid stack in such a manner that they are almost assured of participating in the market. On the other hand, if the prediction interval is wide, they could decide it is too risky to be in the market and possibly incur causer pays fees. Instead they could design the bid stack to be out of the market and instead store the energy in the battery, to be used later.

There is a vast literature on forecasting of solar radiation, over many time scales. Refer to [2] for a critique of the value of deterministic forecasts, that also contains references to several forecasting techniques. Similarly, one of the present authors (Boland), in conjunction with colleagues, described a comparison between using additive deseasoning using Fourier series and autoregressive moving average (ARMA) tools and a separate approach using multiplicative deseasoning with a clear sky model and artificial neural networks (ANN) and also a third using a clear sky model plus ARMA [3]. Significantly, the simpler version of Fourier series plus ARMA performed at least as well over both continental and island locations. There have also been various articles on probabilistic forecasting of solar radiation, using a variety of methods (see [4–7] as examples).

Specifically, in [4], the authors use an ensemble of various activation functions in ANNs to create the prediction intervals. It is applied to datasets obtained from a lab-scale DC micro-grid system for testing. In [5], the systematic change in distribution of the noise terms of the forecasting model is utilised to construct the prediction intervals. This systematic change reflects the fact that the distribution of solar radiation is dependent on the time of day and time of year. For example, the distribution has a higher variance and skewness in the middle of the day in summer than at any other time. In [6], the authors exploit both this systematic change in variance, as well as the conditional change in variance to construct the prediction intervals. The conditional change in variance is typical of a series displaying the ARCH effect, the autoregressive conditional heteroscedastic or variance clustering effect. Periods of high volatility are interspersed with periods of low volatility. The authors of [7] use quantile regression and picking the error distribution to work with based on bins for the errors for varying levels of point forecasts. This is perhaps the most comprehensive of the present methods. Ensemble forecasting of clear sky index followed by post processing of the raw forecasts was used in [8]. Reforecasting using ANN methods is utilised in [9] to improve the skill of three methods, physical based on cloud movement, ARMA and k

nearest neighbour approaches. Seeing that the physical model is probably for forecasting solar radiation, we assume that the others are too. Presumably we can infer from this that an indirect method whereby radiation is forecast and then the plant power output is modelled from that needs further enhancement. From the inspection of the literature it would appear that for solar power, the article of [7] stands out in that they forecast the power output by using the history of power output. In [10], ARMA and generalised autoregressive conditional heteroscedastic (GARCH) models are used to construct prediction intervals for solar radiation. They do not take into account the systematic change in variance that is evident in [5,6]. There are several other articles we could refer to here but we would like to simply point to [11], where the authors use a deep convolutional neural net for their forecasting. We want to point out that while this approach is very useful, we choose to take a simple tool where we can utilise the physical and statistical attributes of the output to build our prediction intervals. One very valuable review article is [12], where the authors describe in some detail a vast array of tools from the simple ARIMA type to ANNs to sophisticated machine learning tools.

The literature is more extensive on probabilistic forecasting of wind. The authors of [13] present both parametric and non-parametric methods for probabilistic forecasting of wind power, but not output from wind farms. Probabilistic forecasting of wind power based on pinball loss optimisation is presented in [14]. Gaussian mixture models are used in [15] for constructing prediction intervals. Ensemble methods have been used from [16] to more recently [17]. Ensemble methods come in different varieties. Ensembles of future trajectories can be developed using physical models of the atmosphere run with different initial conditions. Alternatively, they can be realisations of trajectories constructed using different models, including statistical ones.

In summary, the methods for probabilistic forecasting of wind and solar energy are varied. They range from use of ensemble methods, particularly for wind speed, to parametric methods using assumptions that the error distributions are either Gaussian, Laplace, or other, to non-parametric methods such as quantile regression. In the literature, there is little evidence of direct methods of constructing prediction intervals for solar or wind farms. By direct methods, we mean by using the history of the output from a farm as the basis for forward prediction.

One main contribution of this paper is that we use the power output as the input variable for the forecasting, for both wind and solar farms. One reason that we would advocate this approach is that it is difficult to know how to measure the wind at a wind farm or solar radiation at a solar farm to include the spatial variability inherent over the area of an installation. This approach also enables us to take advantage of inherent physical attributes of the output. The attribute for wind farms that we take into account is that of the so-called ARCH effect, the concept that periods of high volatility are followed by periods of low volatility. In other words, the variance changes conditionally, the variance at time t depends on the variance at previous times. For solar farms, as in solar energy forecasting [5], there is the systematic change in distribution—the distribution of the errors changes over the day. In [5], there a corresponding change over the year. However, for solar farms, the oversizing of the field compared to the inverters means that there is essentially no significant seasonal variation over the year, as on a clear day in winter the output hits the capacity, as well as in the summer.

2. Materials and Methods

The data used in this study come from Snowtown wind farm in South Australia (latitude 33.78° S, longitude 138.21° E), and Broken Hill solar farm in New South Wales (latitude 31.96° S, longitude 141.46° E). The wind farm data are instantaneous power output at five minute intervals from 2017–2018, and the solar farm data are also five minute power output at the same time interval. There are a number of stages of the wind farm at Snowtown and the one we are using has a capacity of 98 MW, while the Broken Hill solar farm has a capacity of 54 MW.

The general procedure is to divide the dataset into training and testing sets. Since for both the solar and wind farms there are two years of data, we have simply taken the first year as the training set and the second as the testing data. The solar farm output has inherent diurnal seasonal variation, while the wind farm output does not. The attributes have both been tested to confirm this using spectral methods.

2.1. Wind Farm Model

Since there is no seasonal component in the wind farm output, we will deal with it first. Note that in this section, we give details of the model constructed in the training phase. The classical time series forecasting approach is via an autoregressive moving average (ARMA(p,q)) model.

$$x_t = \alpha_0 + \alpha_1 x_{t-1} + \alpha_2 x_{t-2} + \dots + \alpha_p x_{t-p} + y_t - \beta_1 y_{t-1} - \beta_2 y_{t-2} - \dots - \beta_q y_{t-q} \quad (1)$$

The α_i, β_i coefficients are estimated and are subject to specific constraints in order to describe a stationary model. In the case of the Snowtown wind farm, detailed analysis results in a simple autoregressive model of order 2 (AR(2)) being the most parsimonious one.

$$x_t = 0.195 + 1.027x_{t-1} - 0.321x_{t-2} + y_t \quad (2)$$

The one ahead forecast thus becomes

$$\hat{x}_t = 0.195 + 1.027x_{t-1} - 0.321x_{t-2} \quad (3)$$

In the simplest of situations for constructing prediction intervals around the forecast, we would have the error y_t display the characteristics of white noise, that is they would be independent and identically distributed (iid). However, while the y_t are uncorrelated, they are dependent. This means that they exhibit an autoregressive conditional heteroscedastic (ARCH) effect. This is evident when one examines the autocorrelation for y_t^2 , the proxy for the variance, as if the variance were uncorrelated there would be at most 5% of the lags being outside the confidence bounds—see Figure 1. The standard method for dealing with this effect would be to use an ARCH or perhaps a generalised ARCH or GARCH model for forecasting the variance. One could then construct appropriate error bounds for the forecast. However, this procedure would assume that the errors are close to normally distributed. As we see from Figure 2, this is not the case (the p -value for a test of normality is <0.005).




Figure 1. The sample autocorrelation function for the error squared for the wind farm.

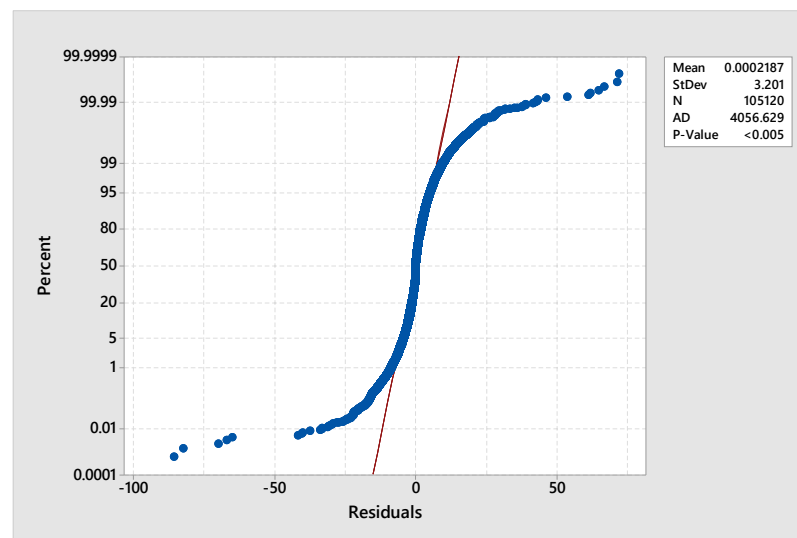


Figure 2. The Q–Q plot to check for normality of the noise for the wind farm.

Thus, we have to transform the errors to the standard normal distribution. To do so, define the combined collection of errors as the random variable Y . Note that in reality there is a random variable for each time t , but for the wind farm we are assuming that the distributions for each time are similar so we can group all the errors together as coming from the same distribution. This assumption follows from the lack of seasonality in the wind farm output. Then for the error at each time t , find the probability $p = P(Y \leq y_t)$. Then, transform each y_t by finding $z_t = F^{-1}(p)$ where $F(z)$ is the cumulative distribution function of the standard normal distribution Z . The usual procedure then is to square the z_t and construct a forecast for the variance using either an ARCH or GARCH model, whichever is appropriate. However, as found in [6], an exponential smoothing forecast model is an efficient alternative. For the estimate of the variance at time t , S_t , we have

$$\begin{aligned} M_t &= \gamma z_{t-1}^2 + (1 - \gamma)M_{t-1}, \\ M_1 &= z_1^2 \end{aligned} \quad (4)$$

The optimal value of γ is selected by minimising the sum of squared differences between the model and the data, as in linear regression. In the case of Snowtown, the forecast is

$$\begin{aligned} M_t &= 0.1493z_{t-1}^2 + (1 - 0.1493)M_{t-1}, \\ M_1 &= z_1^2 \end{aligned} \quad (5)$$

This will give the forecast of the variance, that is it will give an estimate of the variance at time t from the history up to that time. We then construct the upper and lower bounds (in the standard normal space) for the prediction interval by $U_p = z_a \times s_t$, $L_p = -z_a \times s_t$, where $s_t = \sqrt{M_t}$, and z_a is the value of the standard normal distribution for a specified confidence level, for example 1.96 for a 95% confidence level. Subsequently, we calculate $r_u = P(Z < U_p)$ and $r_l = P(Z < L_p)$, and then the bounds for the prediction interval are given by

$$\begin{aligned} U_x &= \hat{x}_t + G^{-1}(r_u) \\ L_x &= \hat{x}_t + G^{-1}(r_l) \end{aligned}$$

where $G(y)$ is the cumulative distribution function of the errors.

2.2. Solar Farm Model

There are inherent differences between the forecast model for the solar farm output versus the wind farm output. The first one is that there is an inherent seasonality in the solar farm output. The first step is to identify and model the form of the seasonality. We have identified several significant cycles using spectral analysis. Fourier series will be used in this step. Define the output at time t by S_t .

$$FS_t = \bar{S} + \sum_{j=1}^k (a_j \cos(2\pi h_j t) + b_j \sin(2\pi h_j t)) \quad (6)$$

where FS_t is the seasonal component, \bar{S} is the mean solar output, and h_j are the significant frequencies in the data.

It is imperative to note that for solar farms in Australia in particular, the seasonality in most instances is different from the seasonality in solar radiation. On a clear day, the solar radiation profile has a distinct peak—see Figure 3. However, in most installations in Australia, the field is oversized compared to the capacity of the inverters. This results in the output displaying a flat maximum over a number of hours on a clear day in summer as in Figure 4, or a smaller number of hours in winter as in Figure 5.




Figure 3. Solar radiation profile.




Figure 4. Solar farm profile—summer.




Figure 5. Solar farm profile—winter.

There are significant modifications to the structure of the Fourier series representation of the radiation, as given in [18]. For solar radiation, one needed the yearly cycle, daily cycle plus the first two harmonics, plus what are termed the beat frequencies. These modulate the amplitude of the daily cycle to suit the time of year. With solar farms being capped, there is no yearly cycle nor beat frequencies as the amplitude does not change over the year.

Once the Fourier series component has been identified, it is subtracted from the data and the deseasoned data are ready to be modelled.

- Form $r_t = S_t - FS_t$, where S_t is the solar farm power output in MW, and FS_t is the Fourier series representation of the seasonal component of the output;
- Check the sample autocorrelation function (SACF) and sample partial autocorrelation function (SPACF) to see what form of an $ARMA(p, q)$ should be used;
- There are two things to note here:
 - The Fourier series and $ARMA(p, q)$ models are estimated on a training set and then tested on a period of time not in the training set;
 - Many people at this stage use much more esoteric means for modelling, like ANN or other machine learning techniques.

The Fourier series fit is given by

$$\begin{aligned}
 FS_t = & 14.79 - 23.39 \cos(2\pi t/288) - 3.87 \sin(2\pi t/288) \\
 & + 10.41 \cos(4\pi t/288) + 3.41 \sin(4\pi t/288)
 \end{aligned} \quad (7)$$

After this is subtracted from the original data, the resulting r_t is fitted with an $ARMA(2, 1)$ model, determined by finding the optimum values of p, q .

$$r_t = 1.241r_{t-1} - 0.257r_{t-2} + b_t - 0.643b_{t-1} \quad (8)$$

We now describe how to construct the prediction intervals for the solar farm output. In this situation, the desire is that b_t are iid, but as with the wind farm noise terms, this is not the case. However, similar to the situation in [5], there is a systematic change in distribution of the noise terms. There is a difference though, because of there being no yearly cycle in the original output data. Therefore, instead of needing to cater for change of distribution for both time of day and time year, here we only need to consider changes due to time of day.

Thus, the algorithm for producing prediction intervals is:

1. Deposit the noise terms in the training set into 24 separate bins for each hour of the day (obviously night-time hours are not needed but this simple method negates the need to determine how day length changes throughout the year). What this means is that at each time step, the noise, or error, is placed into a repository corresponding to the time of day;
2. Depending on what confidence level is desired, determine the appropriate quantiles to correspond to the ends of the intervals, for example for a 95% prediction interval, find the 0.025 and 0.975 percentiles of the errors of the process— L and U ;
3. For time t , get the forecast level $S_t + \hat{r}_t$ and form the ends of the interval;
4. $S_t + \hat{r}_t + L$;
5. $S_t + \hat{r}_t + U$.

2.3. Summary of Methods

There are significant differences between the approaches for wind farms versus solar farms. These mainly result from the fact that the wind farm output studied, common to other wind farms in Australia that we have examined, does not exhibit any significant seasonality. The solar farms we have studied do not show significant yearly seasonality, but obviously display daily seasonality. This basic difference leads directly to the different approaches for the two systems.

2.3.1. Wind Farm

No seasonality model is needed and only a simple autoregressive model is needed for the point forecast. This is estimated using the year of training data. Then, the procedure for developing the prediction intervals for the year of testing data is as follows:

1. Gather the noise terms for the whole year of the training set, forming a probability distribution function (pdf);
2. At each time t in the training set, apply a normalising transform to the noise, using the cumulative distribution function (cdf) of the noise terms;
3. Square the normalised noise terms to form proxies for the variance;
4. Construct an exponential smoothing forecast model for this time series, as in Equation (5);
5. For the testing set, at time t , transform the noise term using the procedure above and the cdf of the training set, and then square the result;
6. Use Equation (5) to forecast the variance for time $t + 1$. Form the lower and upper bounds for the forecast, using the standard deviation thus estimated and the standard normal distribution score corresponding to the level of prediction needed—for example ± 1.96 for a 95% prediction interval;
7. Back transform these values to the noise distribution from the training set;
8. Add these bounds to the point forecast for time $t + 1$ made at time t , to form the desired prediction interval;
9. Repeat for each time in the testing set.

2.3.2. Solar Farms

The approach for the solar farms is slightly more complicated for the point forecast and less complicated for the prediction interval forecast. For the point forecast, we need to model the seasonality, as well as the conditional movement of the mean, as shown in an ARMA model—or a neural network model in other works.

For the prediction intervals, we are considering only the systematic change in distribution. The schema is as follows.

1. If the standard assumptions held, the way to build error bounds around the forecast would be to just to take the standard deviation of the white noise, multiply it by $\pm z$ where the value of z corresponds to the level of probability one wants—e.g., 1.96 for a 95% prediction interval. Then add those values to the forecast;
2. Since the distributions are not normal, one instead finds, for a 95% prediction interval, the 0.025 and 0.975 percentiles of the errors of the process and add them to the forecast;

3. Another complication—the error distributions change over the day. Therefore, we perform this process separately for each hour of the day. For solar radiation, they change over the day and the year.

2.4. Evaluation of Results

Validating the procedure for constructing interval forecasts is different from that of point forecasts. For point forecasts, one is concerned with how close the forecast values are to the subsequently observed values, and that the forecasts are not biased towards systematically over or under forecasting. Thus, the measures that are employed include mean bias error (MBE), root mean square error (RMSE), mean absolute error (MAE), usually both normalised by division by the mean value of the test set, and skill score—see [2].

For evaluating the performance of probabilistic forecasts, one is keen to ensure that the forecasted intervals encompass the data at the required confidence level and that the intervals are not too wide. In other words we test for reliability (or coverage) and sharpness (interval width).

Thus, we go through this procedure:

- In order to validate the approach, we first check the coverage. If you try for a 90% prediction interval, 90% of the observations should be within the intervals [19];
- However, the width of the interval is also important, the smaller the better if also including the observed value within it;
- Note though that there are times when the observed value is just outside the prediction interval;
- Therefore, to take into account this feature, and also the width of the intervals, we look at the interval score for intervals formed for a $(1 - \alpha)100\%$ interval [20].

$$IS_{\alpha} = \frac{1}{N} \sum_{i=1}^N [(U_i - L_i) + \frac{2}{\alpha}(L_i - x_i)1_{x_i < L_i} + \frac{2}{\alpha}(x_i - U_i)1_{x_i > U_i}] \quad (9)$$

Note that in the calculations, this score is normalised by dividing by the mean output of the wind or solar farm.

It is insufficient to simply perform these calculations without comparing the results to some benchmark. Therefore, for solar farm output, we construct a smart persistence forecast. Pick a clear day's output. Get the profile of that day—call it the clear sky model (CSM). There are various ways to do this. One can either search for the clearest day possible in the period with longest day-length in the training set. Or one can be a bit more systematic and select the maximum value for each 5 min period in the whole training set and then smooth this profile using Fourier series. This latter method is what we have employed—see Figure 6. Then, for each t , calculate the ratio of the output at that time versus the clear day profile. For time $t + 1$, find the output that has the same ratio to the CSM. For the wind farm output, we proceed in a simpler manner since there is no seasonality. Note that if one were to be forecasting wind speed rather than wind farm output, there may well be seasonality. In a farm, particularly onshore, this seasonality appears to be smoothed out, due, probably, to the topography and the workings of the turbines. Note also that even with a potential seasonality of wind speeds, this is likely to be intermittent due to calming of the winds from time to time. Therefore, we use simply persistence as the benchmark, $\hat{y}_t = y_{t-1}$.

As well as comparison with a benchmark forecast model, often a comparison is performed with respect to results from other articles in the literature. However, in [21], the authors point out that a score obtained by a forecast method is not a measure of the skill of the forecast since the results depend strongly on the sky conditions of the site. They further quote [22] as stating that the meteorological conditions of the site impact the quality of solar probabilistic forecasts. In conclusion, the score for the reference model at the site reflects the difficulties of forecasting for the site and comparison with that score shows the skill of the proposed model.

Another difficulty of comparisons with the results of other research is that different time frames are used for validation, and also different protocols for construction of measures. An example of the latter is from [23]. Although we, and others, normalise the interval score by dividing by the mean output, they use the difference between the maximum and minimum output. This is not exactly for the interval score but for the PINAW, which is simply the average length of the prediction intervals. In [19], the authors mention that the normalisation can be by the mean irradiation or the installed capacity. Obviously this can be altered slightly, as we have, to mean output.

Therefore, all of these reasons make comparisons with other studies problematic. As a result, we have focused on comparisons with reference forecasting tools. In future work, it is possible to compare with other tools directly using the same data and protocols, but if the model is very complicated that may be difficult unless the code is made available.




Figure 6. Clear sky model.

3. Results

We present the results for both solar and wind farms. The measures that we report are the coverage and interval score. We reiterate that in both cases we train on one year's data and test on another year of data. It must be noted that this includes using the errors from the training set as the basis for transforming the errors of the testing set to normal distributions. One could possibly augment this error set with the errors as we go through the testing set to expand the error set, but we have not done this.

3.1. Wind Farm Results

In Figures 7 and 8, we present the output over two days, first for the present model and also for persistence forecast. There is not a lot of difference apparent, but if one looks closely, the present model has lower peaks and higher troughs in several places. The results for the coverage and interval scores are presented in Table 1. It should be noted here that unlike what we did for the solar farm persistence prediction intervals, we did convert to normal and use that model the variance forecast. We determined that even though this is a partial involvement of our modelling procedure, it was justified since the errors could be treated as a whole, rather than depositing them in separate bins. This could well explain how the interval scores are closer for the wind farm than the solar farm. One could surmise that if we had simply taken the naive approach for the persistence prediction intervals for wind farm output, the interval scores would be farther apart. Even though the interval scores for persistence are very good, the coverage is not sufficient.

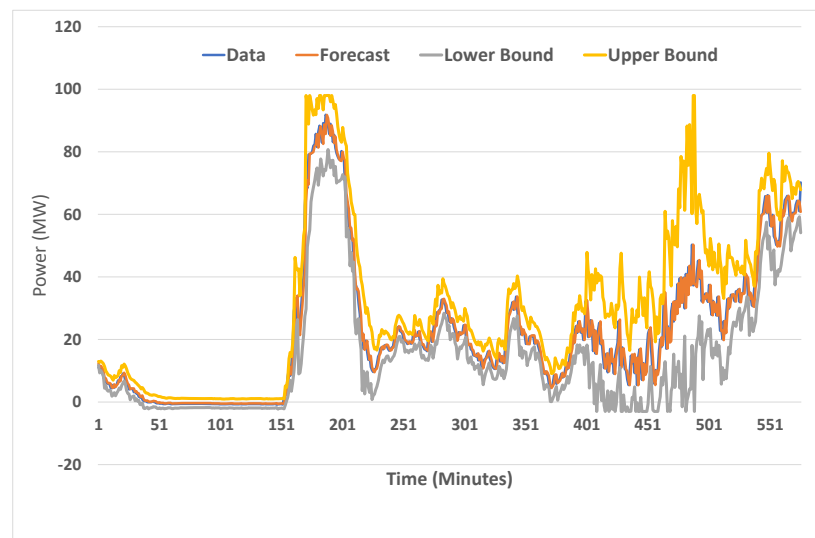


Figure 7. Probabilistic forecasting example for Snowtown wind farm.

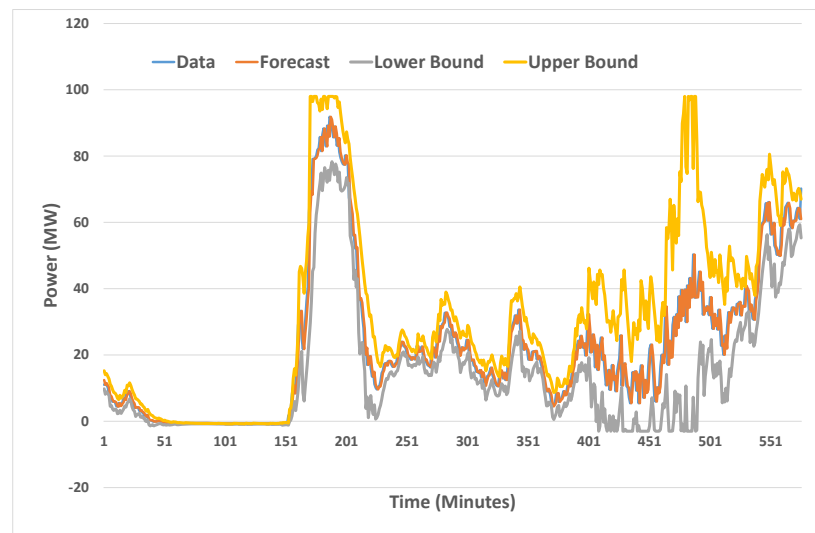


Figure 8. Probabilistic forecasting example for Snowtown solar farm (Persistence).

Table 1. Wind farm coverage and interval score for different confidence levels—comparison with persistence.

Measure	Confidence Level—Percentage	Model	Persistence
Coverage	90	93.2	85.4
		Interval Score	0.246
Coverage	95	97.6	91.4
		Interval Score	0.348

3.2. Solar Farm Results

In Figures 9 and 10, we present the output for the present model and smart persistence, over two days, one that has intermittent generation and one subject to mostly clear sky. To construct the prediction intervals for smart persistence, we simply added the $\alpha/2$ and $1 - \alpha/2$ quantiles to the forecasts, where $(1 - \alpha) \times 100$ is the required confidence level. We decided that if we used the bins for the errors and performed transformations to normal, it would be too much like adding a feature of our modelling procedure to the smart persistence forecast.

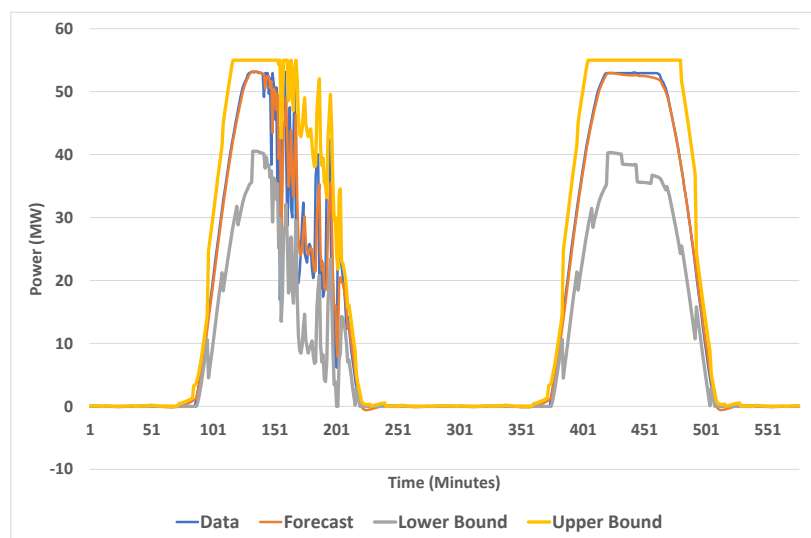


Figure 9. Probabilistic forecasting example for Broken Hill solar farm.

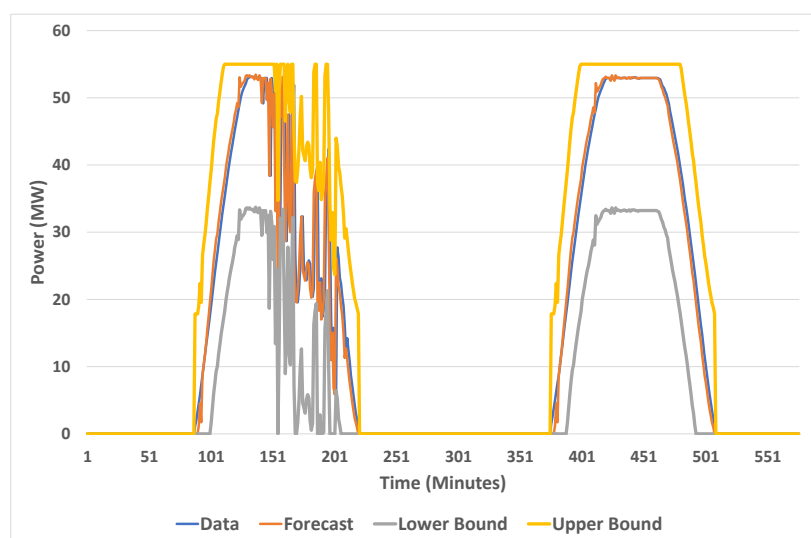


Figure 10. Probabilistic forecasting example for Broken Hill solar farm with smart persistence.

In Table 2, we can see the comparison between the present model and smart persistence in terms of coverage and interval score. The coverage is much closer to the target for the present model, while also having significantly lower interval score.

Table 2. Solar farm coverage and interval score for different confidence levels—comparison with smart persistence.

Measure	Confidence Level—Percentage	Model	Smart Persistence
Coverage	80	80	88.7
		Interval Score	0.478
Coverage	90	90.6	96.1
		Interval Score	0.658
Coverage	95	95.2	98.5
		Interval Score	0.716

4. Discussion

One can take either of two approaches when trying to forecast output from wind or solar farms. One can either forecast the resource, wind speed or solar radiation, and use a power conversion model (PCM) to obtain the forecast for the output. Or one can use the historical output and forecast directly from that. We favour the latter approach for several reasons.

- To forecast the resource and then use a PCM introduces two levels of model error, rather than one;
- The PCM model can become complicated when one is not directly converting to power, but to power output that is, as shown here, artificially capped in the case of solar farms by oversizing the field. This may also be already happening with wind farms in Australia, but not with the one studied here. There are arguments for so doing as detailed in [24];
- There is the more philosophical question—perhaps actually technical—of where do you measure the resource at a wind or solar farm? A large scale wind farm, as is the case in Australia, is spread over a fairly large and usually undulating terrain. If one measures the wind speed at one location, this may well be quite different from what wind impacts turbines at another part of the farm. In the case of a solar farm, a cloud may be over the sensor but not over other parts of the farm, or vice versa.

We have shown how one can effectively construct prediction interval forecasts for situations where the common assumptions for time series forecasting—that the errors are independent and identically distributed—do not hold. There are other methods to do so like quantile regression, but we have chosen a relatively simple approach. We tested the methods versus the benchmark of smart persistence in the case of solar farms, and simple persistence in the case of wind farms, with good results. The methods can easily be applied to any wind or solar farm. One proviso is that if the solar farm does not have capping due to oversizing the field or some other mechanism, it might be necessary to add more frequencies in the Fourier series seasonality representation.

Future Work

There are a few questions that come to mind. One arises from consideration of the forecasting of solar radiation in [18], where as well as the systematic change in variance of the errors being taken into account, the conditional change (ARCH effect) was catered for, resulting in narrower prediction intervals. This approach will be taken on in subsequent work. The other main question arising is what exactly is a persistence interval forecast in situations where the variance of the errors changes over time? For solar farms, we took a simple approach since we would have had to in essence use some of the techniques developed for our forecast model in the persistence forecast. For the wind farm, we modified the simple persistence approach somewhat since it was easy to perform, but perhaps we actually enhanced the performance of the interval persistence forecast by doing so. We intend to investigate this question further. Another avenue for future investigation is to see if adding exogenous variables can enhance the performance of the model. Examples could include ambient temperature, solar radiation measurements if we can access them, and perhaps wind speed data as well.

Author Contributions: Conceptualization, J.B. and S.F.; methodology, J.B. and S.F.; software, J.B. and S.F.; validation, J.B.; formal analysis, J.B.; investigation, J.B. and S.F.; data curation, J.B. and S.F.; writing—original draft preparation, J.B.; writing—review and editing, J.B. and S.F.; visualization, J.B.; supervision, J.B.; project administration, J.B. Both authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable for studies not involving humans or animals.

Informed Consent Statement: Not applicable for studies not involving human.

Data Availability Statement: The data for this study are available from the authors.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

BOM	Australian Bureau of Meteorology
ARENA	Australian Renewable Energy Agency
AEMO	Australian Energy Market Operator

References

1. Snell, I.M.C.T.; Consani, S.; West, C.S.; Amos, M. *Solar Power Ensemble Forecaster Final Report—Public*; Australian Renewable Energy Agency (ARENA): Canberra, Australia, 2021.
2. Yang, D.; Alessandrini, S.; Antonanzas, J.; Antonanzas-Torres, F.; Badescu, V.; Beyer, H.G.; Blaga, R.; Boland, J.; Bright, J.M.; Coimbra, C.F.; et al. Verification of deterministic solar forecasts. *Sol. Energy* **2020**, *210*, 20–37. [[CrossRef](#)]
3. Boland, J.; David, M.; Lauret, P. Short term solar radiation forecasting: Island versus continental sites. *Energy* **2016**, *113*, 186–192. [[CrossRef](#)]
4. Ni, Q.; Zhuang, S.; Sheng, H.; Kang, G.; Xiao, J. An ensemble prediction intervals approach for short-term PV power forecasting. *Sol. Energy* **2017**, *155*, 1072–1083. [[CrossRef](#)]
5. Grantham, A.; Gel, Y.; Boland, J. Nonparametric short-term probabilistic forecasting for solar radiation. *Sol. Energy* **2016**, *133*, 465–475. [[CrossRef](#)]
6. Boland, J.; Grantham, A. Nonparametric Conditional Heteroscedastic Hourly Probabilistic Forecasting of Solar Radiation. *J* **2018**, *1*, 174–191. [[CrossRef](#)]
7. Golestaneh, F.; Pinson, P.; Gooi, H.B. Very short-term nonparametric probabilistic forecasting of renewable energy generation—With application to solar energy. *IEEE Trans. Power Syst.* **2016**, *31*, 3850–3863. [[CrossRef](#)]
8. Yagli, G.M.; Yang, D.; Srinivasan, D. Ensemble solar forecasting using data-driven models with probabilistic post-processing through GAMLSS. *Sol. Energy* **2020**, *208*, 612–622. [[CrossRef](#)]
9. Chu, Y.; Urquhart, B.; Gohari, S.M.; Pedro, H.T.; Kleissl, J.; Coimbra, C.F. Short-term reforecasting of power output from a 48 MWe solar PV plant. *Sol. Energy* **2015**, *112*, 68–77. [[CrossRef](#)]
10. David, M.; Ramahatana, F.; Trombe, P.J.; Lauret, P. Probabilistic forecasting of the solar irradiance with recursive ARMA and GARCH models. *Sol. Energy* **2016**, *133*, 55–72. [[CrossRef](#)]
11. Wang, H.; Yi, H.; Peng, J.; Wang, G.; Liu, Y.; Jiang, H.; Liu, W. Deterministic and probabilistic forecasting of photovoltaic power based on deep convolutional neural network. *Energy Convers. Manag.* **2017**, *153*, 409–422. [[CrossRef](#)]
12. Ahmed, R.; Sreeram, V.; Mishra, Y.; Arif, M.D. A review and evaluation of the state-of-the-art in PV solar power forecasting: Techniques and optimization. *Renew. Sustain. Energy Rev.* **2020**, *124*, 109792. [[CrossRef](#)]
13. Zhang, Y.; Wang, J.; Wang, X. Review on probabilistic forecasting of wind power generation. *Renew. Sustain. Energy Rev.* **2014**, *32*, 255–270. [[CrossRef](#)]
14. Sun, M.; Feng, C.; Chartan, E.K.; Hodge, B.M.; Zhang, J. A two-step short-term probabilistic wind forecasting methodology based on predictive distribution optimization. *Appl. Energy* **2019**, *238*, 1497–1505. [[CrossRef](#)]
15. Jin, H.; Shi, L.; Chen, X.; Qian, B.; Yang, B.; Jin, H. Probabilistic wind power forecasting using selective ensemble of finite mixture Gaussian process regression models. *Renew. Energy* **2021**, *174*, 1–18. [[CrossRef](#)]
16. Pinson, P. Very-short-term probabilistic forecasting of wind power with generalized logit-normal distributions. *J. R. Stat. Soc. Ser. C Appl. Stat.* **2012**, *61*, 555–576. [[CrossRef](#)]
17. Kim, Y.; Hur, J. An ensemble forecasting model of wind power outputs based on improved statistical approaches. *Energies* **2020**, *13*, 1071. [[CrossRef](#)]
18. Boland, J. Time series modelling of solar radiation. In *Modeling Solar Radiation at the Earth's Surface: Recent Advances*; Springer: Berlin/Heidelberg, Germany, 2008; pp. 283–312. [[CrossRef](#)]
19. Lauret, P.; David, M.; Pinson, P. Verification of solar irradiance probabilistic forecasts. *Sol. Energy* **2019**, *194*, 254–271. [[CrossRef](#)]
20. Gneiting, T.; Raftery, A.E. Strictly proper scoring rules, prediction, and estimation. *J. Am. Stat. Assoc.* **2007**, *102*, 359–378. [[CrossRef](#)]
21. Le Gal La Salle, J.; David, M.; Lauret, P. A new climatology reference model to benchmark probabilistic solar forecasts. *Sol. Energy* **2021**, *223*, 398–414. [[CrossRef](#)]
22. Alessandrini, S.; Delle Monache, L.; Sperati, S.; Cervone, G. An analog ensemble for short-term probabilistic solar power forecast. *Appl. Energy* **2015**, *157*, 95–110. [[CrossRef](#)]
23. Van der Meer, D.W.; Widén, J.; Munkhammar, J. Review on probabilistic forecasting of photovoltaic power production and electricity consumption. *Renew. Sustain. Energy Rev.* **2018**, *81*, 1484–1512. [[CrossRef](#)]
24. McInerney, C.M.; Bunn, D.W. Optimal Oversizing of Wind Generation Facilities. *SSRN Electron. J.* **2015**, 1–41. doi:10.2139/ssrn.2576307. [[CrossRef](#)]