

Article

Constructing Interval Forecasts for Solar and Wind Energy Using Quantile Regression, ARCH and Exponential Smoothing Methods

John Boland 

Industrial AI Research Centre, University of South Australia, Mawson Lakes Boulevard,
Mawson Lakes, SA 5095, Australia; john.boland@unisa.edu.au; Tel.: +61-415051569

Abstract: The research reported in this article focuses on a comparison of two different approaches to setting error bounds, or prediction intervals, on short-term forecasting of solar irradiation as well as solar and wind farm output. Short term in this instance relates to the time scales applicable in the Australian National Electricity Market (NEM), which operates on a five-minute basis throughout the year. The Australian Energy Market Operator (AEMO) has decided in recent years that, as well as point forecasts of energy, it is advantageous for planning purposes to have error bounds on those forecasts. We use quantile regression as one of the techniques to construct the bounds. This procedure is compared to a method of forecasting the conditional variance by use of either ARCH/GARCH or exponential smoothing, whichever is more appropriate for the specific application. The noise terms for these techniques must undergo a normalising transformation before their application. It seems that, for certain applications, quantile regression performs better, and the other technique for some other applications.

Keywords: probabilistic forecasting; prediction intervals; ARCH/GARCH; exponential smoothing; solar irradiation; solar farm; wind farm



Citation: Boland, J. Constructing Interval Forecasts for Solar and Wind Energy Using Quantile Regression, ARCH and Exponential Smoothing Methods. *Energies* **2024**, *17*, 3240. <https://doi.org/10.3390/en17133240>

Academic Editor: Oscar Barambones

Received: 3 April 2024

Revised: 11 June 2024

Accepted: 27 June 2024

Published: 1 July 2024



Copyright: © 2024 by the author. 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

Australia has seen a rapid development of wind and solar farms connected to the National Electricity Market (NEM) in the last two decades. At first, the development mainly focused on wind farms, but the last decade has seen a rapid increase in the number and scale of solar farms. In the NEM, there are three types of generators: scheduled, semi-scheduled and non-scheduled. The scheduled fleet consists of traditional coal- and gas-fired plants, whose operating characteristics are well-known and are less dependent on the weather. Semi-scheduled generators are wind and solar farms over 30 MW and also batteries over 5 MW. Any generators under 30 MW are non-scheduled. Scheduled generators submit a bid stack every five minutes, stating the amount of energy that they will supply in the following five minutes at ten price bands from $-\$1000$ per MWh to $+\$16,000$ per MWh (in \$AUS). If they do not supply what they predict (outside some tolerance level), they are penalised. These act as price makers. Semi-scheduled generators are price takers, but they still are expected to supply a forecast of their output, and, once again, if the forecast is outside some tolerance limit, they are penalised. They can also be required to curtail their output if there is oversupply in the grid. The non-scheduled generators have no restrictions and are simply price takers. The present author was part of a team that helped five solar farms to improve their five-minute forecasting, but that was simply for the expected output, or what can be termed point forecasting—see [1].

In this paper, putting error bounds on the forecast (in other words, constructing prediction intervals) is the focus. Two methods for producing such intervals are canvassed, one using quantile regression. A reason for using this method is that the noise terms in the forecasting method do not follow the hoped-for characteristics of being independent

and identically distributed random variables. The other method is somewhat complicated, and also takes this conundrum into account. In this method, one applies a normalising transformation to the noise terms, then applies tools to forecast their conditional variance, setting error bounds and back-transforming these bounds. The efficacy of these methods is evaluated with a comparison to a naive model. The goal of the research is to ascertain if one of the two methods is superior and in which situations.

In general, the published research has focused on the construction of prediction intervals for either solar resource or wind speed. In Australia, particularly in South Australia, there has been a rapid expansion of renewable-energy-generated electricity. For example, the South Australia Department of Energy and Mining reports that the percentage of electricity generated from renewable sources is, at present, almost 70% (Department of Energy and Mining website, accessed 4 April 2024). The same website states that the goal is 100% by 2030, but a recent government statement suggests it will be closer to 2027 that this happens. It is necessary to consider how setting prediction intervals for forecasts for output from farms differs from setting them for the variable renewable energy (VRE) itself. This is particularly the case for solar farms in Australia, as most have a specific characteristic. For various reasons, they have oversized fields of panels compared to the capacity of the inverters. Figure 1a,b shows that, on clear days in winter as well as summer, Broken Hill solar farm reaches capacity output. This is a common occurrence for solar farms in Australia. A common approach to characterise the seasonality in solar irradiation is by use of a clear sky model and, subsequently, dividing the incoming global solar irradiation by that model to obtain a clear sky index (CSI). This approach does not work with capped output. One would have to devise a power conversion model or construct an empirical clear sky output model.




Figure 1. Clear days in (a) winter and (b) summer.

It is important to find the most skillful statistical methods for constructing prediction intervals, given that, for the very short-term forecasting needed to suit the NEM in Australia, statistical methods are very well-suited. They often perform equally as well as more sophisticated methods—see [2] for a comparison of the performance of statistical and artificial neural network methods for point forecasting.

This paper is set out as follows. Section 2 reviews the relevant literature for interval forecasting and discusses some areas for improvement. This is not specifically for improving the skill of forecasting, but rather for defining methods that are more appropriate for the applications studied here. Section 3 examines the specific methods that are used in this study. Section 4 describes the metrics used for evaluation of the skill of forecasting the intervals, while Section 5 gives the results of applying those metrics. Section 6 gives the concluding remarks.

2. Literature Review

There are various methods for constructing prediction intervals for forecasts of solar energy. David et al. [3] compared various methods for intraday probabilistic forecasting of solar energy. A previous article by David et al. [4] used recursive ARMA and GARCH models for the same task. In that article, they also canvassed many other methods, not

only for the prediction intervals, but also for deterministic forecasting. These included the coupled autoregressive and dynamical systems (CARDS) method for deterministic forecasting, as in [5,6]. The reference model for comparison for prediction intervals is the analog ensemble mode [7]. An important consideration, as the authors point out, is that the use of the GARCH technique for forecasting the variance relies on the assumption that the errors are normally distributed. They go on to point out that, for the two example locations they used for their analysis, the assumption of normality is not supported. A way around this conundrum was explored in [8] using a normalising transformation, and this is discussed further in Section 3. David et al. [3] also pointed to the sieve bootstrap method for generating prediction intervals, and pointed specifically as well to this method being used in a more sophisticated manner by Grantham et al. [6]. In that article, the authors took into consideration the dependence of the noise terms—remaining after the deterministic forecast—on the time of year and the time of day. When constructing the prediction intervals, sampling is carried out solely from the noise distribution for the specific values of the two variables reflecting those dependencies, the sun elevation and hour angle. One other approach for constructing the prediction intervals is the use of quantile regression (see, for example, [9]). In this method, instead of forecasting the mean, as in standard methods of time-series modelling, one forecasts the various quantiles. The method is used on the noise terms and is fully described in Section 3.

Recently, some researchers have been applying methods from the artificial intelligence area to construct prediction intervals. Pan et al. [10] used gated recurrent unit (GRU) neural networks to perform the deterministic forecast for solar irradiation and then added the prediction intervals using kernel density estimation. Based on the forecast skill metric, the GRU method performs well for both 5 and 15 min ahead forecasts. From inspection of the error metrics, the method appears to work well compared with naive approaches. The metrics used are the prediction interval coverage probability (PICP) and prediction interval normalised average width (PINAW). In general, the PICP is higher than desired, and this seems to be a common feature of many procedures. Sun et al. [11] used a machine learning technique called quantile random forest (QRF) to form their prediction intervals for ultra-short-term wind power. As they stated, the random forest algorithm provides forecasts of the mean, but QRF provides a distributional forecast. They compared their method with a number of other candidate procedures, and the PINAW values were superior for all confidence levels tried: 80, 90, and 95. In their analysis, they selected the highest value of PICP for each confidence level as the best, rather than the one closest to the selected level. Alcantara et al. [12] described the use of deep neural networks (DNN) to construct prediction intervals. They investigated the differences between using DNN tools to construct quantiles and then to find prediction intervals and that of directly finding the intervals. They argued that, since the interval width is part of the optimising approach, the direct method should result in narrower intervals. They compared two DNN methods: directly finding the intervals and indirectly finding them using the DNN quantile approach. One of the direct methods performs best over most of the levels tested, with the DNN quantile best at the 95% level. The advantage for the DNN tools in this paper over many other methods is that the PICP is, in most cases, closer to the prescribed level than in the other two papers canvassed. The treatment in the present paper is mainly for separate applications from those that have been canvassed in the literature. While one resource data set, for solar irradiation, is treated in this paper, the other two data sets are for solar and wind farm output, both with capacity limits. As explained in Section 1, the solar farm in particular has characteristics that differ from those that are canvassed in the surveyed literature. This necessitates novel approaches.

Many approaches to both deterministic and probabilistic forecasting of solar energy had their beginnings in the wind energy forecasting sphere. This is probably because large-scale wind energy utilisation in the form of wind farms started before the advent of solar farms of any magnitude. One of the early contributions to the probabilistic forecasting

of wind power—specifically, for the horizon of 48 to 72 h ahead—was that in [13]. In more recent times, machine learning tools have been used for similar purposes [14,15].

3. Methods

3.1. Data

There are three data sets used in this paper. One data set, that of solar irradiation, is for Cabauw, Netherlands (51.96 N, 4.90 E) and is 15 min data in kilowatt hours (kWh). The second data set is solar farm output data at a 5 min time scale for Broken Hill solar farm (31.98 S, 141.39 E), NSW, Australia. The third data set is wind farm output on a half-hour time scale for Clements Gap wind farm (33.48 S, 130.10 E), South Australia, Australia. The solar and wind farm data are power in megawatts (MW). All training is carried out for one year's complete data set and testing is carried out on a separate year's data.

3.2. Deterministic Forecasting

There are several methods for performing the deterministic forecasting of solar energy. The first step is to model the seasonality of the data in the training set. Many researchers, including David et al. [4], deal with seasonality by constructing the clear sky index, which is the global horizontal irradiation divided by a suitable clear sky model. They used the McClear clear sky model [16]. After that, many researchers used artificial neural nets (ANNs) to forecast the clear sky index. David et al. [4] used the autoregressive moving average (ARMA) method in recursive mode. A different approach to determining the seasonality is followed here, by constructing the expected value of the solar energy at any time t using Fourier series. In [17], there is a description of the physical nature of the significant frequencies that are inherent in the solar radiation data. The determination of the significant frequencies is performed using spectral analysis, identifying the frequencies that contribute most significantly to explaining the variance in the series. The yearly and daily cycles are intuitively obvious. The necessity of including the twice daily cycle, also identified by this method, is not obvious. This could represent the fact that, as well as night being different from day, morning is different from afternoon. This is because the sky is usually more turbid in the afternoon. The frequencies just surrounding the daily and twice daily cycles, at 364, 366, as well as at 729, 731 cycles per year, are included. These beat frequencies, also called sidebands, are well-known in signal processing. In the language of that discipline, one can have a carrier signal with frequency $\omega_c = 2\pi f_c$ that has its amplitude modulated by a signal at lower frequency $\omega_m = 2\pi f_m$. In the specific case examined here, f_c is the daily cycle, while f_m is the yearly cycle. What occurs is that the daily amplitude is modulated by the yearly cycle. If the sidebands are not included, the Fourier series representation exhibits peculiar behaviour, including anomalous positive solar energy at night in summer (see Figure 2 for summer in Cabauw, Netherlands as an example) and negative at night in winter. Alternatively, if one includes the sidebands for the same days, we can see that the Fourier series gives a sensible representation at night—see Figure 3. Note that the amplitude is such that it also gives a more sensible peak to the Fourier series. As a result, the seasonality model takes the form

$$\begin{aligned}
 F_t = & \alpha_0 + \alpha_1 \cos \frac{2\pi 365}{35040} t + \beta_1 \sin \frac{2\pi 365}{35040} t + \\
 & \sum_{n=1}^2 \sum_{m=-1}^1 (\alpha_{nm} \cos \frac{2\pi(365n+m)t}{35040} + \\
 & \beta_{nm} \sin \frac{2\pi(365n+m)t}{35040}).
 \end{aligned} \tag{1}$$




Figure 2. Effect of removing the contribution of sidebands.




Figure 3. Effect of including the contribution of sidebands.

The next step is to take the difference between the original data S_t and the seasonal model to form the residual series $R_t = S_t - F_t$. By analysing the sample autocorrelation function (SACF) and sample partial autocorrelation function (SPACF) (see [18] for details of the procedure), we find that the best forecast model for the residuals is given by an autoregressive model with 3 lags $AR(3)$, so that the forecast at time $(t - 1)$ for time t , \hat{R}_t is given by

$$\hat{R}_t = \gamma_1 R_{t-1} + \gamma_2 R_{t-2} + \gamma_3 R_{t-3}. \quad (2)$$

When the $AR(3)$ forecast is added to the Fourier series, this results in the final forecast. Figure 4 shows the one-step-ahead forecast as an example.

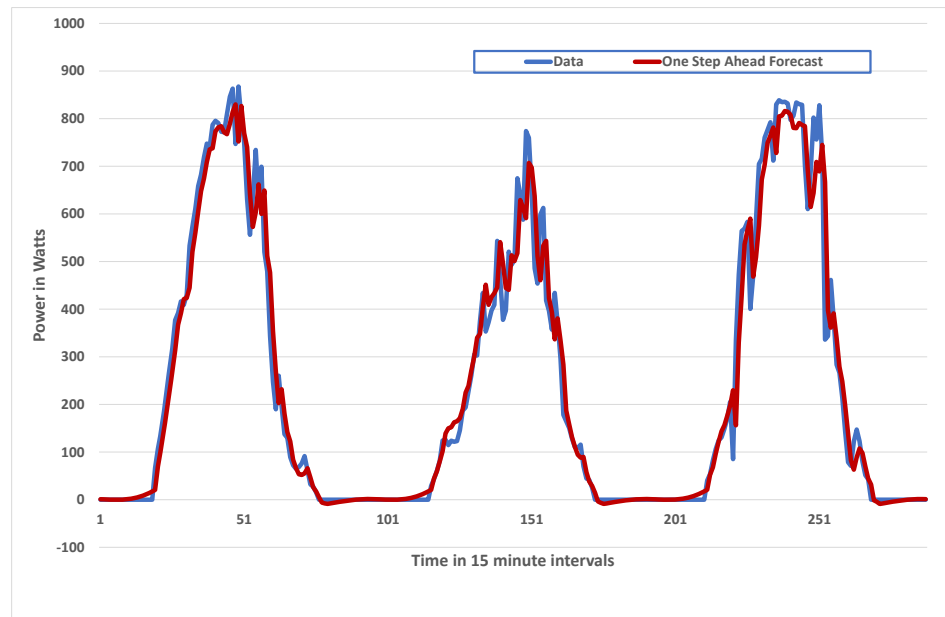


Figure 4. One-step-ahead forecast overlaid on the data.

3.3. Quantile Regression

Any one-step-ahead statistical forecasting method can be encapsulated by the structure

$$Y_t = f(F; Y_{t-1}, \dots, Y_{t-p}; X_{i,t-1}, \dots, X_{i,t-q}) + Z_t \quad (3)$$

This contains the seasonality F and any autoregressive qualities, plus any connection with exogenous variables, if applicable. Knowledge of the statistical qualities of Z_t is necessary in order to construct the error bounds of the forecast. In this formulation, it is hoped, and sometimes assumed, that Z_t is independent and identically distributed (i.i.d.). Further, it would make the construction of the prediction intervals straightforward if the distribution were Gaussian. But, for solar irradiation, none of these assumptions hold. To estimate the error bounds, or the limits of the prediction intervals, one method is quantile regression. This is performed on Z_t , the noise, which is derived by subtracting the final forecast model from the data but filtered so that we are working only on data for which the solar elevation is greater than 10° . Quantile regression does not require any assumptions about the noise distribution.

For quantile level τ of the response, the goal is to

$$\min_{\beta_0(\tau), \beta_1(\tau), \dots, \beta_p(\tau)} \sum_{i=1}^n \rho_\tau(y_i - \beta_0(\tau) - \sum_{j=1}^p z_{ij} \beta_j(\tau))^2. \quad (4)$$

$$\rho = \tau \max(r, 0) + (1 - \tau) \max(-r, 0) \quad (5)$$

is the check function. If the error in the regression in a single period, r , is positive, then the check function multiplies the error by τ (and by $1 - \tau$ if negative). In this study, the predictor variables are the previous five lagged values of the noise. Note that, when performing the optimisation for the one-step-ahead forecast for time $t + 1$ at time t , we are regressing Z_{t+1} on $Z_t, Z_{t-1}, \dots, Z_{t-5}$.

3.4. ARCH or GARCH Variance Forecast

In [8], the method that was used for generating prediction intervals managed to avoid the problem of having to assume that the noise terms were normally distributed before being able to use ARCH or GARCH methods. First, it is instructive to define these methods. The noise Z_t remaining after the deterministic forecast has essentially no autocorrelation, so the usual assumption from that is that the Z_t are independent. But, as first identified

in financial time series, often the Z_t^2 series has autocorrelation. This means that the series displays an ARCH effect, and, since Z_t^2 is a proxy for the variance in the original series, the variance changes as t changes. That is the origin of the term **AutoRegressive Conditional Heteroscedastic** or ARCH. If there is a sudden break in significant lags for the partial autocorrelation function for Z_t^2 , this series can be modelled as a purely autoregressive, or $AR(p)$, process. This ARCH model is given by Equation (6):

$$\sigma_t^2 = \alpha_1 Z_{t-1}^2 + \alpha_2 Z_{t-2}^2 + \dots + \alpha_p Z_{t-p}^2 \quad (6)$$

When the sample autocorrelation function dies down very slowly and the sample partial autocorrelation function does also, but with a slightly less long tail, then this is typical of what drew Bollerslev [19] to expand on the work of Engel [20], who developed the ARCH model by generalising it to the generalised ARCH (GARCH) model. What this formulation achieves is replace a long-run autoregressive model for the variance with a much more parsimonious autoregressive moving average (ARMA) model—albeit with a slight alteration.

The $GARCH(m, s)$ model is of the following form:

$$\begin{aligned} Z_t &= \sigma_t \epsilon_t \\ \sigma_t^2 &= \alpha_0 + \sum_{i=1}^m \alpha_i Z_{t-i}^2 + \sum_{j=1}^s \beta_j \sigma_{t-j}^2 \end{aligned} \quad (7)$$

Here, $\epsilon_t \sim (0, 1)$ and one obtains estimates of the α_i and β_j in the following manner.

Take $\{Z_t^2\}$ as an observed series and fit an $ARMA(p, q)$ model to this series, with parameter estimates $\hat{\phi}_i$ and $\hat{\theta}_i$. Then,

$$\begin{aligned} \hat{\beta}_i &= \hat{\theta}_i \\ \hat{\alpha}_i &= \hat{\phi}_i - \hat{\theta}_i \end{aligned} \quad (8)$$

As mentioned in Section 2, use of the ARCH or GARCH approach assumes that the noise is at least close to normally distributed. An inspection of the noise from the deterministic forecasting for 15 min data from Cabauw, Netherlands with a normal curve with the same mean and standard deviation shows that it is definitely not normal—see Figure 5. This finding is in line with that of David et al. [4]. The standard result from this form of forecasting is noise that has a Laplace distribution, leptokurtic with fat tails. This feature will not be used here.

Simply applying an ARCH or GARCH forecast to the squared noise is not appropriate. Instead, first transform the noise to the equivalent value from the standard normal distribution. Figure 5 shows the noise for all values of solar elevation above 10° as an illustration. The reality is that there are different distributions for different hours of the day, reflecting the changing volatility over the day, and this change in distribution must be accounted for. The day is split into the hours before 6 a.m. each hour between 6 a.m. and 6 p.m., and the hours after 6 p.m. The algorithm for setting the prediction intervals is given in Algorithm 1. Figures 6 and 7 illustrate the transform procedure. Take a value of the noise z_t , and, using the cumulative distribution function (CDF) for the noise in Figure 6, find the probability value τ for which $F(z_t, i) = \tau$. In terms of the figure, locate z_t on the horizontal axis, then progress vertically to the curve. Then, move left horizontally to the vertical axis, and where one lands is the value of τ . Then, Figure 7 shows how we take τ and find the corresponding value from the standard normal distribution. In essence, find τ on the y -axis, move horizontally to the normal CDF and drop down to obtain the corresponding normal distribution value.

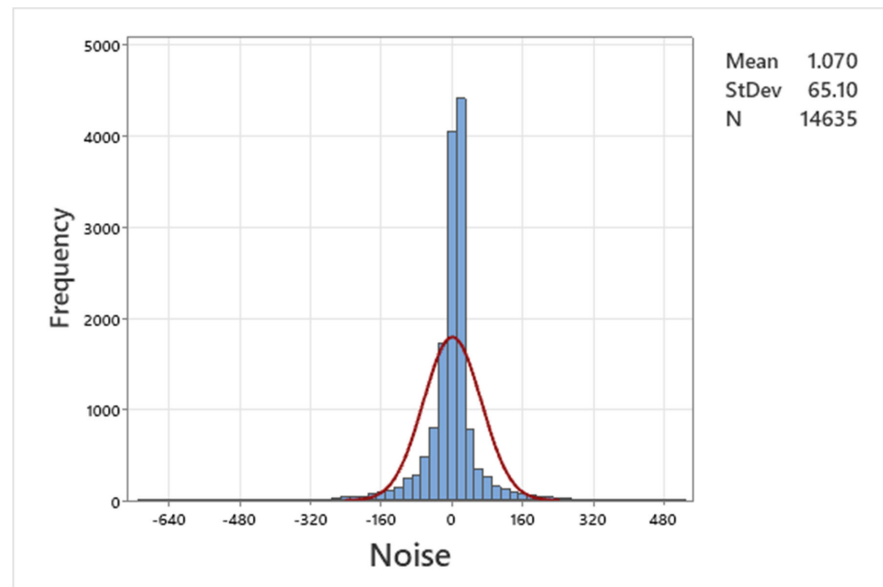


Figure 5. Distribution of the noise from deterministic forecasting. The blue denotes the distribution, and the red curve the normal distribution with the same mean and variance.




Figure 6. Finding the probability of occurrence of z_t .

Inspiration for this novel approach comes from an article produced early on in the research into forecasting solar radiation [21]. In that article, Sfeir transformed the solar radiation data to normal before applying an ARMA forecasting model. They assumed that ARMA was not an appropriate tool unless the data were at least close to normally distributed. For simple forecasting, this assumption is not supported, but, for other applications like synthetic generation, it is a sensible approach. However, a simple transformation of the data in their entirety is not sufficient. As is seen in the algorithm, there are differing distributions for different times of the day, and this aspect must be taken into account.




Figure 7. Finding the equivalent value of the standard normal distribution.

Algorithm 1: The Transformation to Normal Algorithm

Data: The noise terms from the deterministic forecasting

for $t = 1, \dots, N$ **do**

 Find $F(Z_t, i)$, the cumulative probability distribution (CDF) of Z_t for time of day i ;

 Transform each specific value z_t according to $\gamma_t = F^{-1}(z_t, i)$, with $\gamma_t \sim N(0, 1)$;

 Form γ_t^2 for every 15 min over the year;

 Find the GARCH forecast model for γ_t^2 ;

 Use the GARCH model to estimate the variance for each 15 min, and then calculate the standard deviation σ_t ;

 Form the lower and upper bounds for the prediction interval (PI), for instance for a 95% PI, they are $\pm 1.96 \times \sigma_t$;

 Find the corresponding probabilities from the cumulative distribution function of the standard normal distribution p_L, p_U ;

 Find the values of $F(Z_t, i)$ corresponding to p_L, p_U , using the suitable CDF in each case; **Result:** Add these bounds to the deterministic forecast, to form the final prediction intervals.

end

3.5. Exponential Smoothing Variance Forecast

The algorithm described in Section 3.4 is used in the same way for forecasting the variance, simply substituting an exponential smoothing forecast tool instead of ARCH or GARCH. This is introduced since, for some situations, this approach is more appropriate. To describe the process, write the forecast model for σ_t^2 —see Equation (9):

$$\hat{\sigma}_t^2 = \alpha \gamma_{t-1}^2 + (1 - \alpha) \sigma_{t-1}^2 \quad (9)$$

3.6. Wind Farm Modelling

The model construction for setting prediction intervals for wind farm output is much simpler than for solar energy. This is due to two particular facts. One is that there is very little or no seasonality. Perhaps wind speeds may display some seasonality but, largely, wind farm output in the Australian context does not. The farms are large and

extend mainly over either a very large landscape or a hilly terrain near coastal areas. Both of these attributes tend to smooth out any seasonality of the wind itself. Thus, the deterministic forecast is simply an ARMA model. The other factor is that, even though the noise distribution from the forecast model is not normal, it does not vary over the day or year. The transformation to normal errors can be performed with a single CDF of the noise. The determination of the prediction intervals is performed following the algorithm in Section 3.4, but with simply a single-noise distribution.

3.7. Solar Farm Modelling

The procedure for determining a one-step-ahead forecast and then putting error bounds on the forecast for solar farm output is similar to that used for solar irradiance. There is one exception to this. In Australia, most of the solar farms, including the Broken Hill one, have an oversized field of panels compared to what the inverters can handle. This results in a capping of the output, as seen in Figure 8, so we can say that the capacity of the solar farm is approximately 54 MW. The capacity is reached on clear days throughout the year, resulting in the only significant frequencies in the data being the once a day and twice a day. There is no yearly cycle of note and, thus, the beat frequencies are also redundant. As for the solar irradiance forecasting, once the seasonality that is modelled with Fourier series is removed, an $AR(p)$ process is sufficient to complete the point forecasting, as demonstrated in Figure 9. The prediction intervals are found using both the quantile regression and normal transformation methods as with the solar irradiance.

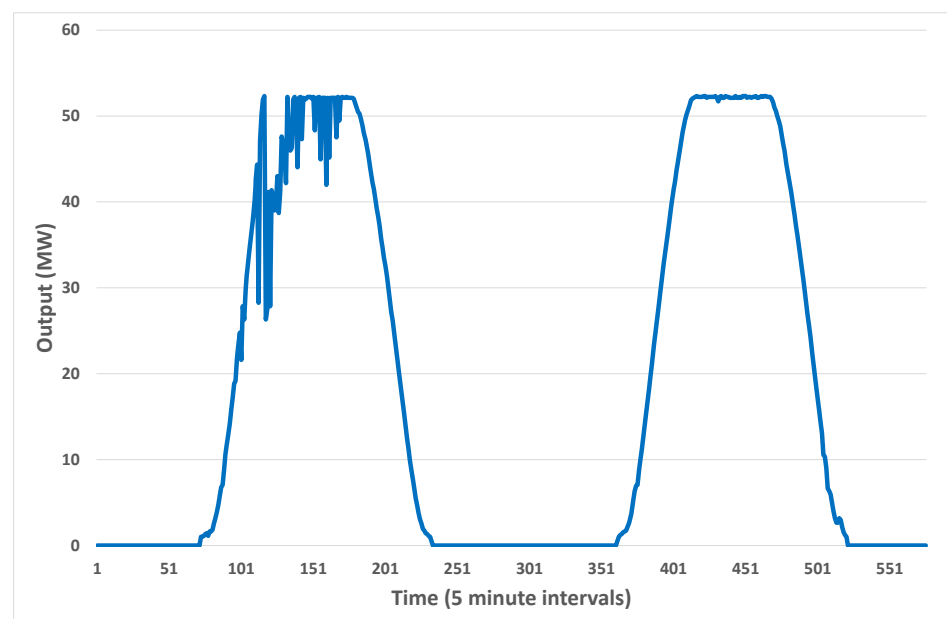


Figure 8. Two days of Broken Hill solar farm output.

3.8. Benchmark

The goal is to compare the performance of the quantile regression approach with using the normalising transformation plus ARCH or GARCH. It is useful as well to demonstrate how they compare to a fairly simple naive method. It is only applied to the prediction interval estimation, as it still involves using the deterministic forecast described in Section 3.2. A completely naive approach for the deterministic forecast would entail using a persistence forecast for the expected output $X_{t+1} = X_t$. Note that, often, researchers will only use the persistence forecast on the data once the seasonality is removed and term it smart persistence, but this negates the fact that the form of seasonality determination is in itself a model. The naive approach here is to take the deterministic forecast using the seasonality and ARMA model, and then add to and subtract from it the standard deviation of the total noise distribution (for solar elevation above 10°) multiplied by the score from the standard

normal distribution corresponding to the probability level desired. This is solely for the solar energy evaluation. Since there is little to no seasonality for the wind farm output, we use the simplest benchmark: a persistence forecast with normal errors around that forecast.

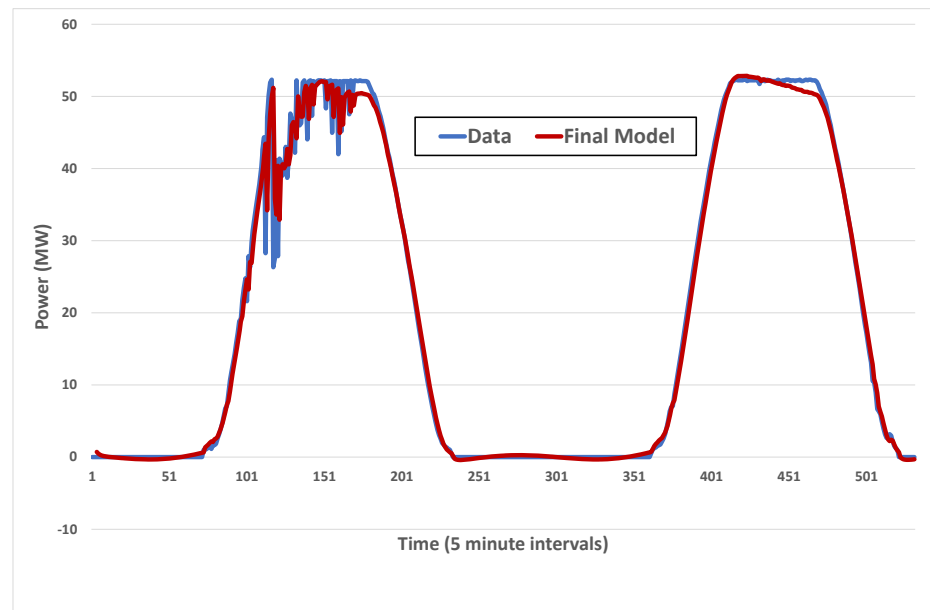


Figure 9. Two days of Broken Hill solar farm output and forecast.

4. Metrics

Two simple metrics are used for comparison. This is because the goal is to compare the two methods, quantile regression and ARCH/GARCH, for constructing the prediction intervals, and also to compare these methods with the benchmark. This is performed for solar energy, solar farm output and wind farm output. The first metric is coverage. When a 95% prediction interval is calculated, for example, it is expected that close to 95% of the observed values will fall within the intervals. There are two comments to make about this stipulation. Some researchers would state it somewhat differently, that at least 95% of the observed values would fall in the intervals. Overcoverage is also deficient, as it implies that, potentially, the intervals are wider than necessary. This leads into to the second point. The other criterion is sharpness. The coverage criterion should be maintained with the narrowest intervals possible. Thus, the mean interval width is also used for evaluation.

The metric for calculating the coverage is the prediction interval percentage coverage (PICP):

$$PICP = \frac{1}{n} \sum_{i=1}^n \delta_i \quad (10)$$

In this equation, δ_i is a Boolean operator, equal to unity when the observed value S_t falls within the prediction interval for time t , and zero when it does not.

For the interval width, the metric is the prediction interval normalised average width (PINAW):

$$PINAW = \frac{1}{nR} \sum_{i=1}^n \eta_i \quad (11)$$

In this equation, η_i is the width of the i th prediction interval and R is the difference between the minimum and maximum for solar energy or wind farm output (or the capacity for solar farm output). Why it is not the capacity for the wind farm is that, for wind farms, there are intervals for which the output is negative. This is because the control systems remain on even if there is no wind, resulting in power being drawn from the grid, resulting

in negative output. The periods of zero wind can happen randomly. There is not the same problem with the solar farm output since sunrise and sunset can be calculated for all days of the year and, between these times, the output is always positive. The control systems activate and deactivate at the specific times of sunrise and sunset.

5. Results and Discussion

As well as reporting the metrics, some graphical results are displayed to give a pictorial version of the comparison. The most telling results come, however, from the quantitative metrics.

5.1. Solar Energy

First, a pictorial summary of the comparison of the three methods is given—quantile regression, an ARCH model and the naive approach that simply assumes that the noise terms are identically distributed normal variables. Figure 10 shows some interesting results for the 95% probability level. For example, it appears that, at least at this probability level, the ARCH model performs much better than the other two on clear days and worse on intermittent days. The other two models perform overall about the same. On the other hand, Figure 11 shows that the ARCH model performs significantly better than the naive model for the 80% probability level for clear days and at least as good for intermittent days. In this figure, LB stands for lower bound and UB for upper bound.




Figure 10. Comparing prediction intervals using quantile regression and an ARCH forecast with the naive approach.

A more explicit way to compare the methods is by use of the two metrics, PICP and PINAW. From inspection of Table 1, for PICP, it is obvious that the quantile method gives the most precise results. On the other hand, the method that stands out as being deficient according to this metric is the one where the noise terms are considered to be independent and identically distributed normal random variables. Table 2 gives the normalised mean widths of the prediction intervals for the three methods. Combining the results for both metrics, the most reliable method is the model where the noise terms are transformed, with respect to the systematic changes of variance, to the standard normal distribution. The transformed noise values are then squared and an ARCH model is used to forecast the variance. Then, lower and upper bounds are found and the values are then back-transformed.




Figure 11. Comparing prediction intervals using an ARCH forecast with the naive approach.

Table 1. PICP for the methods for solar irradiation.

	Naive	Quantile	ARCH
99	96.3	98.9	98.8
95	93.7	94.9	96.0
90	91.8	90.0	92.4
80	88.4	79.9	83.3

Table 2. PINAW percentages for the methods for solar irradiation.

	Naive	Quantile	ARCH
99	31.9	47.4	33.9
95	24.6	27.1	23.2
90	20.7	18.1	17.3
80	16.3	9.4	11.0

5.2. Wind Farm Output

From Figure 12, it is very hard to tell which works best. Examination of the metrics in Tables 3 and 4 is more revealing. If one simply looks at the interval widths, then it appears that the naive model performs better at the 95% level. The coverage at that level is 93.9%, which would mean a narrower width on average than if the coverage were more in keeping with the desired coverage. It does not perform at all well in comparison at the other levels. The two approaches trialled here appear to behave very similarly, with perhaps a slight advantage to the quantile regression method. The added benefit of that method is that it is simpler to use, so, at least for the wind farm studied here, the quantile regression method is preferred.

Table 3. PICP for Clements Gap wind farm.

	Naive	Quantile	ARCH
95	93.9	95	96.0
90	91.0	90	91.2
80	85.5	80	80.7

Table 4. PINAW percentages for Clements Gap wind farm.

	Naive	Quantile	ARCH
95	31.1	33.1	33.8
90	27.5	24.8	25.5
80	21.8	17.2	17.2




Figure 12. Comparing prediction intervals for Clements Gap wind farm.

5.3. Solar Farm Output

In Figure 13, we compare the use of quantile regression with the transform method for 95% prediction intervals. In the figure, Q stands for quantile and T for transform, with L for lower bound and U for upper. The interesting result is that the quantile regression method appears to perform better when the data are highly variable, producing narrower bounds, while the transform method seems to be better on the clear days. If one examines the coverage in Table 5, it is obvious that the naive approach is not performing well, with the coverage for the nominal 80% level being 91.0%. And, in Table 6, apart from the single instance of a narrow mean interval at the 95% level, once again the naive approach is not performing well. If one takes the overall set of results for the quantile regression versus transform methods, there is not a lot of difference. Overall, with much narrower means intervals for the 95% and 90% levels, the suggestion is that the transform method is slightly better.

Table 5. PICP for Broken Hill solar farm.

	Naive	Quantile	Transform
95	94.9	96.2	94.0
90	93.4	92.2	89.8
80	91.0	83.6	78.1

Table 6. PINAW percentages for Broken Hill solar farm.

	Naive	Quantile	Transform
95	29.6	34.4	26.3
90	25.4	22.8	20.6
80	20.3	11.1	13.3




Figure 13. Comparing prediction intervals for Broken Hill solar farm.

The various methods were evaluated for three applications: solar irradiation and solar and wind farm output. The first observation is that both the quantile regression and transform methods outperformed the naive method in general. The exception is for high probability levels but, in most of these cases, the interval width was better, but the coverage was not as good. At these very short-term time scales, very simple methods work well, so this is a positive result.

For solar irradiation, if one examines Figure 10, the conclusion could be that the quantile method performs better for highly variable data, while the transform method performs better on clear days. When examining the overall metrics, they point to the transform method being superior in the main. For the wind farm output, there is not any obvious difference between the two approaches. For the solar farm output, examination of the metrics points to the transformation method as the better of the two. An interesting result from the pictorial depiction is that the transform method appears to perform better with variable data, while, on clear days, the quantile approach performs better. This is opposite to what happened with the irradiation data. Looking at these results, perhaps a mixture of the two methods might be the best option. An attempt to do so by simply taking the mean of the two options did not add significantly to the skill. From that, future work should include the use of machine learning techniques to blend the approaches.

It is not a simple matter to compare this results of the modelling here with that of what has appeared in the literature. Since the Australian context is different from the rest of the world, both in the short time scale of the electricity market and how solar farms are managed—with the fields often oversized—there can be only moderate comparisons made without applying the methods directly. One point to make is that, apart from [12], the other methods do not, in general, have very good PICP results, while, apart from the naive method, the methods in this study do have that feature. It is difficult to compare PINAW results when the PICP values do not match.

6. Conclusions

This paper gives a comparison of two approaches to construction of error bounds for short-term forecasting for solar irradiation as well as solar and wind farm output. In the context of the short-term operation of the electricity grid and market, the generation of prediction intervals will be crucial in decision making with increased use of storage mechanisms linked to renewable energy installations. This paper is not about forecasting of solar and wind farm output per se, but rather about putting error bounds around those

forecasts. It is pointed out in the literature review that quantile regression and its variants have been used for constructing such bounds, and one of the methods used here is quantile regression. The other method is not one that is routinely present in the literature. Yes, ARCH and GARCH have been used for this purpose, but not with the added feature of first transforming the noise terms to Gaussian before applying the conditional forecasting of the variance. By performing the normalising transformation, then forecasting the variance and back-transforming the bounds, the problem of non-normality of the noise is overcome. This approach in general outperforms the quantile regression method. The overall conclusion from this study is that the transform method is superior in performance, even though the quantile regression method is simpler to implement.

Funding: This research received no external funding.

Data Availability Statement: The original contributions presented in the study are included in the article, further inquiries can be directed to the author.

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

References

1. Snell, T.; West, S.; Amos, M.; Farah, S.; Boland, J.; Kay, M.; Prasad, A. *Solar Power Ensemble Forecaster Final Report Project Summary and Findings*; Technical Report; CSIRO: Canberra, Australia, 2020.
2. Boland, J.; David, M.; Lauret, P. Short term solar radiation forecasting: Island versus continental sites. *Energy* **2016**, *113*, 186–192. [[CrossRef](#)]
3. David, M.; Luis, M.A.; Lauret, P. Comparison of intraday probabilistic forecasting of solar irradiance using only endogenous data. *Int. J. Forecast.* **2018**, *34*, 529–547. [[CrossRef](#)]
4. 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](#)]
5. Huang, J.; Korolkiewicz, M.; Agrawal, M.; Boland, J. Forecasting solar radiation on an hourly time scale using a Coupled AutoRegressive and Dynamical System (CARDS) model. *Sol. Energy* **2013**, *87*, 136–149. [[CrossRef](#)]
6. Grantham, A.; Gel, Y.; Boland, J. Nonparametric short-term probabilistic forecasting for solar radiation. *Sol. Energy* **2016**, *133*, 465–475. [[CrossRef](#)]
7. Alessandrini, S.; Monache, L.D.; Sperati, S.; Cervone, G. An analog ensemble for short-term probabilistic solar power forecast. *Appl. Energy* **2015**, *157*, 95–110. [[CrossRef](#)]
8. Boland, J.; Grantham, A. Nonparametric Conditional Heteroscedastic Hourly Probabilistic Forecasting of Solar Radiation. *J* **2018**, *1*, 174–191. [[CrossRef](#)]
9. Lauret, P.; David, M.; Pedro, H.T. Probabilistic solar forecasting using quantile regression models. *Energies* **2017**, *10*, 1591. [[CrossRef](#)]
10. Pan, C.; Tan, J.; Feng, D. Prediction intervals estimation of solar generation based on gated recurrent unit and kernel density estimation. *Neurocomputing* **2021**, *453*, 552–562. [[CrossRef](#)]
11. Sun, Y.; Huang, Y.; Yang, M. Ultra-Short-Term Wind Power Interval Prediction Based on Fluctuating Process Partitioning and Quantile Regression Forest. *Front. Energy Res.* **2022**, *10*, 867719. [[CrossRef](#)]
12. Alcántara, A.; Galván, I.M.; Aler, R. Direct estimation of prediction intervals for solar and wind regional energy forecasting with deep neural networks. *Eng. Appl. Artif. Intell.* **2022**, *114*, 105128. [[CrossRef](#)]
13. Pinson, P.; Nielsen, H.A.; Møller, J.K.; Madsen, H.; Kariniotakis, G.N. Non-parametric probabilistic forecasts of wind power: Required properties and evaluation. *Wind Energy* **2007**, *10*, 497–516. [[CrossRef](#)]
14. Wang, J.; Wang, S.; Zeng, B.; Lu, H. A novel ensemble probabilistic forecasting system for uncertainty in wind speed. *Appl. Energy* **2022**, *313*, 118796. [[CrossRef](#)]
15. Wan, C.; Cui, W. Machine Learning-Based Probabilistic Forecasting: A Combined Bootstrap and Cumulant Method. *IEEE Trans. Power Syst.* **2023**, *39*, 1370–1383. [[CrossRef](#)]
16. Lefèvre, M.; Oumbe, A.; Blanc, P.; Espinar, B.; Gschwind, B.; Qu, Z.; Wald, L.; Schroedter-Homscheidt, M.; Hoyer-Klick, C.; Arola, A.; et al. McClear: A new model estimating downwelling solar radiation at ground level in clear-sky conditions. *Atmos. Meas. Tech.* **2013**, *6*, 2403–2418. [[CrossRef](#)]
17. Boland, J. Characterising seasonality of solar radiation and solar farm output. *Energies* **2020**, *13*, 471. [[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. Bollerslev, T. Generalized Autoregressive Conditional Heteroscedasticity. *J. Econom.* **1986**, *31*, 307–327. [[CrossRef](#)]

20. Engle, R.F. Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica* **1982**, *50*, 987–1007. [[CrossRef](#)]
21. Sfeir, A.A. A stochastic model for predicting solar system performance. *Solar Energy* **1980**, *25*, 149–154. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.