



MERIDIAN ENERGY AUSTRALIA: Wind Forecasting Demonstration Project LESSONS LEARNT REPORT 6

Project Details

Recipient Name	Meridian Energy Australia
Primary Contact Name	Ben Coffey
Contact Email	Ben.Coffey@infrastructurecapital.com.au
Reporting Period	September 2021 – February 2022

This Project received funding from ARENA as part of ARENA's Advancing Renewables Program. The views expressed herein are not necessarily the views of the Australian Government, and the Australian Government does not accept responsibility for any information or advice contained herein.

EXECUTIVE SUMMARY

This report outlines lessons that Meridian Energy Australia and the University of Melbourne have recently learnt regarding clustered wind turbine-based forecasting and wind turbine cross correlations for the Mt. Mercer wind farm.

KEY LEARNINGS

Lesson learnt No.1: Breaking the Mt. Mercer wind farm into larger number of sub-wind-farm-clusters did not yield an increase in accuracy to the forecast beyond 7 clusters.

Category: Technical

Objective: Demonstrate the ability to submit five-minute ahead self-forecasts via AEMO's web based MP5F API.

Details: Our *a priori* hypothesis for this project was that sub-diving the wind farm into groups of similar producing wind turbines and summing a forecast for each group would yield an overall more accurate forecast than a single power forecast. This hypothesis was based on the assumption that our errors, provided they were “truly random”, would cancel each other out when aggregating and our forecast accuracy would increase.

The trade-off for this gain in accuracy, would be the increased difficulty in forecasting (namely requiring 64 data input streams for each turbine versus the single power data input stream). However, in an offline forecasting scenario (where live data availability was not an issue), our forecast error decreased with an increasing number of clusters only to a point, as shown in Figure 1. Here our “Ensemble” forecast in red is a combination of lightGBM and xgboost, whereas our lightGBM forecast in yellow is simply the lightGBM forecast. These two machine learning approaches were selected for two reasons, namely they were robust (able to handle missing values) and they performed well in preliminary tests. After 18 clusters, the forecast accuracy did not further increase. Note that our Figure only shows to 20 clusters for readability, but this trend continued to the point where our analysis was terminated.

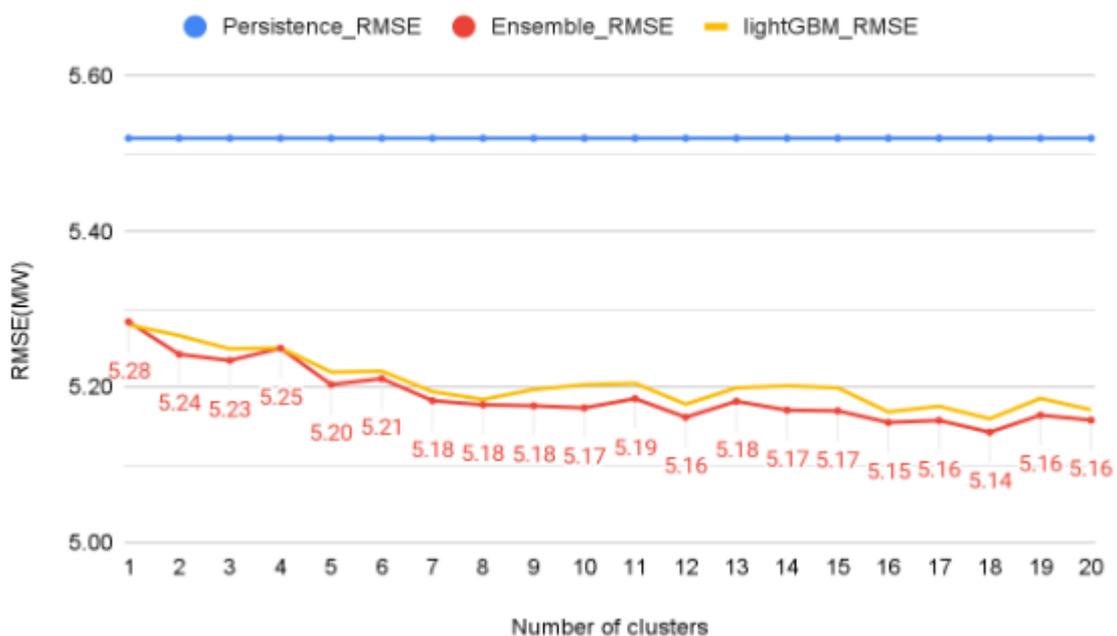


Figure 1: Forecast accuracy for increasing number of wind-turbine clusters for the Mt. Mercer wind farm.

The surprising part of this analysis was the modest level of increase for a larger increase in complexity. That is, despite needing to increase the number of forecasts from 1 to 18, we only decreased RMSE from 5.28 MW to 5.14 MW, and MAE from 3.14 to 3.08 MW, respectively.

We do stress, however, that this analysis was performed offline. It is uncertain how well this type of forecast would perform in a live, real-time setting.

Implications for future projects: The results from this analysis suggest that each forecaster should weigh the benefits of sub-cluster forecasts for themselves. That is, there may be some accuracy improvements to be gained (~1.9%) for the loss of simplicity of the forecast. Our forecast team has favoured simplicity for live forecasting, but other forecasters may prefer to eliminate as much error as possible.

Lesson learnt No.2: There was a noticeable increase in the ability to train machine-learning based forecasts for month-based data splits relative to naïve data-chunk splits.

Category: Technical

Objective: Demonstrate the ability to submit five-minute ahead self-forecasts via AEMO’s web based MP5F API.

Details:

In our cluster analysis, we also determined that there was a noticeable change in forecast skill depending on our data split. Our results from Figure 1 were for our “ideal” split, where we split the data on a per month basis, following Figure 2.

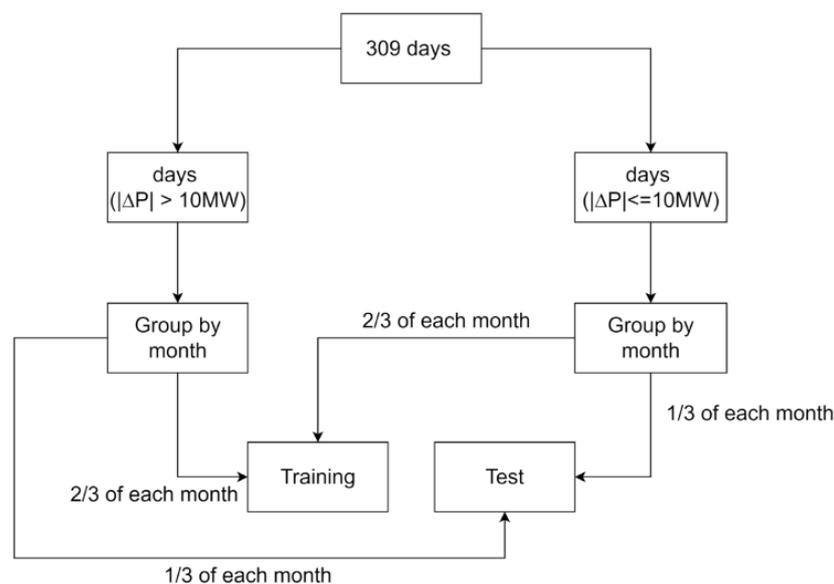


Figure 2: Ideal data split for machine-learning based models for Mt. Mercer wind farm. Here ΔP refers to the total change in wind farm power over the 370 second forecast horizon.

Doing so noticeably reduced our forecasting errors relative to the “naïve” split where the Training data was July-2019 through Feb-2020 and the Test data was Mar-2020 through June-2020.

This can be seen in Figure 3, where the same analysis as Figure 1 was repeated, just with this new (naïve) train-test data split was employed.

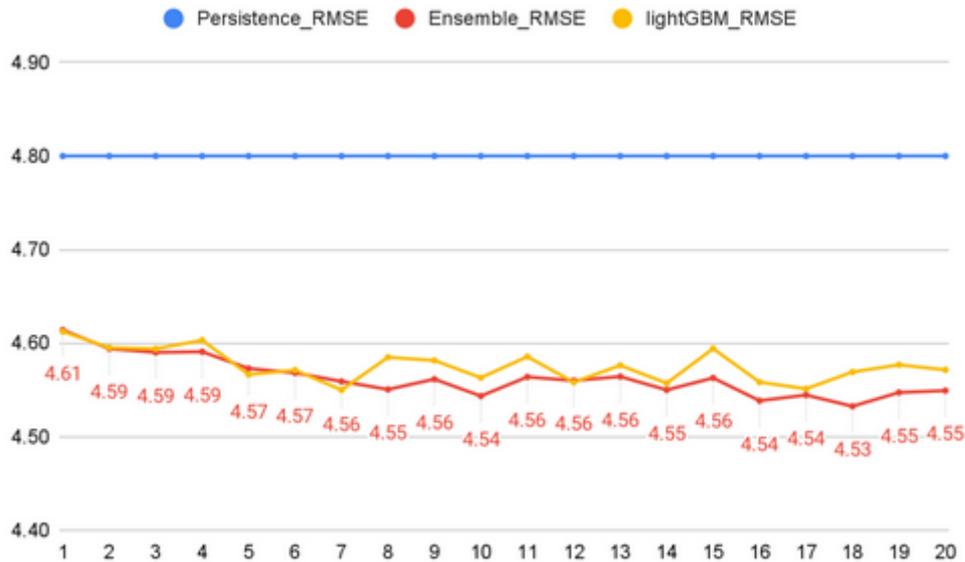


Figure 3: Forecast accuracy for increasing number of wind-turbine clusters for the Mt. Mercer wind farm with naïve data split.

While the absolute errors seem smaller for the naïve split (where the error based on the persistence forecast is also smaller due to the months that we are testing), the relative performance increase for 18 clusters in Figures 1 and 3 correspond to a 6.85% and a 5.57% increase over persistence, respectively. Herein we use persistence error to denote a forecast that employs the value from the previous timestep for the forecast. This is the simplest type of forecast, and the one that works best for a random process. However, the outperformance of persistence by these forecasts suggests that the process is not truly random, and hence is forecastable. The fact that the seasonal split outperformed the naïve split, suggests that this forecastability can be improved by accounting for season-related wind changes.

Implications for future projects: Projects seeking to employ machine learning for their forecasts should be aware of the season changes of the wind. Namely, a traditional naïve training-testing split of the data may learn certain weather patterns better than others. By instead sub-dividing each month, the model is exposed to many seasonal patterns in training, leading to a better overall forecast.

Lesson learnt No.3: Cross correlation values for power generation of wind turbines remain high for only very small durations and distances.

Category: Technical

Objective: Other (analysis of wind power data)

Details: Using our June 2019 – June 2020 data from the wind farm, we sought to determine if wind direction could somehow be used to increase forecasting ability. In order to do this, we began to examine the cross correlation (in short, how well two time-series correlate in general and the delay between the two signals that leads to optimal matching of signals).

We began this analysis by looking at closely spaced wind turbines for periods of time where the wind was especially stable (remaining within 90 degrees of north for the plots herein). By closely spaced, we mean that the turbines are spaced no more than 1 km from one another.

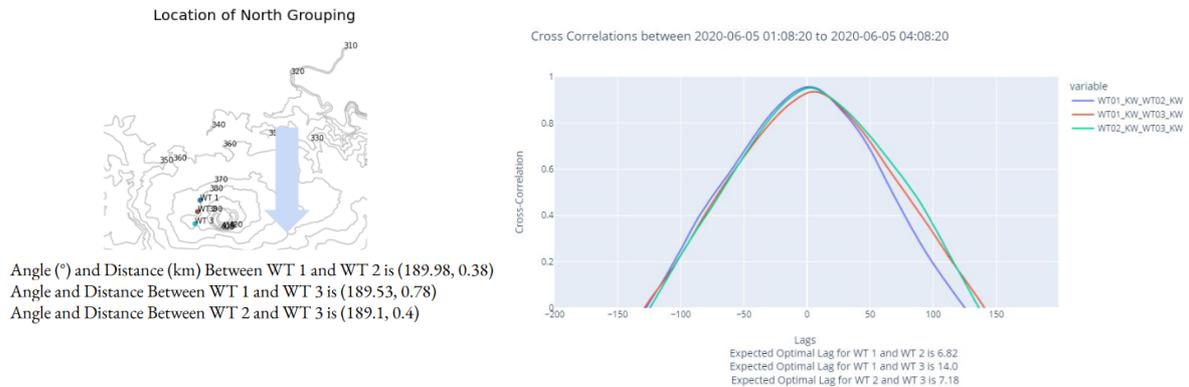


Figure 4: Cross correlation of close wind turbines when the wind was out of the North. Note that lag values correspond to 10-second intervals.

Figure 4 shows an expected cross-correlation behaviour, the plots peak at approximately the expected value, and the power values correlate well for the period of time that Taylor’s Frozen Turbulence Hypothesis yields, before decaying towards zero. Also note that the furthest two turbines (WT1 and WT3) have the lowest peak value and widest width of lag values where the cross correlation remains positive. Other periods of time had somewhat similar behaviour (though depending on the wind it was possible for the cross correlation value to remain positive from -200 to 200).

However, when we considered wider distances of greater than 1 km, the cross correlation plots could become illogical. This is possibly best seen in Figure 5, where optimal coherence was often in negative time, despite being defined such that all optimal coherence values should be positive.

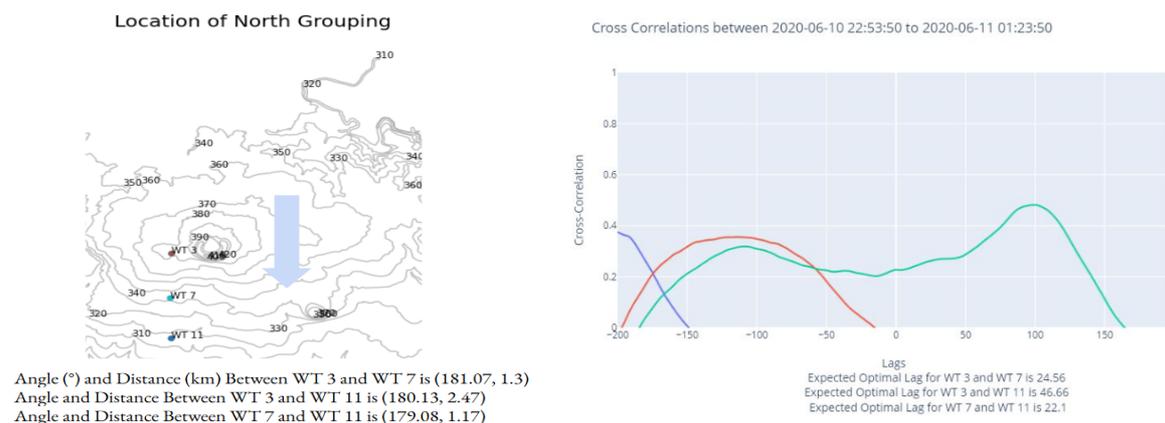


Figure 5: Cross correlation of moderately spaced wind turbines when the wind was out of the North. Note that lag values correspond to 10-second intervals.

We note that some of the wider distance cross correlation plots behaved well like those of the close plots, but the breakdown with increased distance, does suggest a possible barrier

of increased forecasting accuracy when trying to account for wind propagating through the farm.

Implications for future projects: Projects should be very careful how they implement wind data in their forecasts. That is, as the wind propagates through the farm, it does not appear to behave following Taylors Frozen Turbulence hypothesis. Instead, the wind continues to evolve and hence the power generation values become uncorrelated with increasing distance.