Wind & Solar Forecasting for the NEM

17 February 2021

PALISADE

MONASH University

Advisian
Worley Group

advisian.com

Table of contents

Table list

Figure list

Acknowledgements

This project received funding from the Australian Renewable Energy Agency (ARENA) as part of ARENA's Advancing Renewables Program.

We would also like to acknowledge Christoph Bergmeir from Monash University and Josh Lowndes of Palisade Partners for their significant contributions and support to the project.

Disclaimer

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.

Acronyms and abbreviations

Acronym/abbreviation	Definition
AEMO	Australian Energy Market Operator
ANEMOS	Proprietary Forecasting Model Used by AEMO
API	Application Programming Interface
ARIMA	Autoregressive Integrated Moving Average Model
ASEFS	Australian Solar Energy Forecasting System
AWEFS	Australian Wind Energy Forecasting System
CPF	Causer Pays Factor
ETS	Exponential Smoothing Model
FCAS	Frequency Control Ancillary Services
ITB	Inverter Transformer Block
MAE	Mean Absolute Error
NEM	National Energy Market
RMSE	Root Mean Square Error
RNN	Recurrent Neural Network Model
RRSF	Ross River Solar Farm
SCADA	Supervisory Control and Data Acquisition
SDF	Semi Dispatch Flag
WWF	Waterloo Wind Farm

1. Project summary

Advisian is a participant in the ARENA funding initiative for improving 5-minute ahead self-forecasting for wind and solar farms operating in the National Electricity Market (NEM). Advisian partnered with Monash University and Palisade to develop the forecasting models. Both Models are machine learning models which draw on SCADA data as inputs to the model.

Two semi-scheduled generators were chosen for the study. The Waterloo Wind Farm (WWF) in South Australia, comprising thirty-seven V-90 and six V-117 Vestas Wind Turbines with a total installed capacity of 130.8 MW, and the 116 MW Ross River Solar Farm (RRSF) in Queensland.

1.1 Waterloo Wind Farm data set

The raw data used to train the model comprises extracts from the SCADA system and other sources in the wind farm. This includes alarm and maintenance data, constraints data, wind speed, wind power, wind direction and the availability of turbines, all in spreadsheet formats, spanning from January 2018 to April 2020. Total generated wind power and AWEFS forecast was also downloaded and processed the NEMWeb website. Whilst every wind farm is obliged to report to AEMO data comprising active power, wind speed, wind direction, and ambient temperature with values at least every 4 seconds, Waterloo wind farm only historised data as minutely averages, which is what was used for this project.

The wind model does not require additional sensors such wind masts, meteorological stations or lidar, which greatly simplifies the implementation of the solution. The operational model uses windspeed and active power only. In addition, AEMO Semi Dispatch Flag (SDF), possible power and wind turbine availability inputs are used to train the model.

1.2 Ross River Solar Farm data set

Data used to develop the forecasting models contain each inverter transformer block's (ITB) power production, various measures of solar irradiance and temperature for the period February 2019 to January 2020 at 1-minute intervals. Two months of this data set (December 2019 to January 2020) was set aside for testing. In addition, periods in the data set which contained many gaps (constraints) were excluded.

Figure 1-1 Solar farm training & testing months split by constrained and unconstrained data

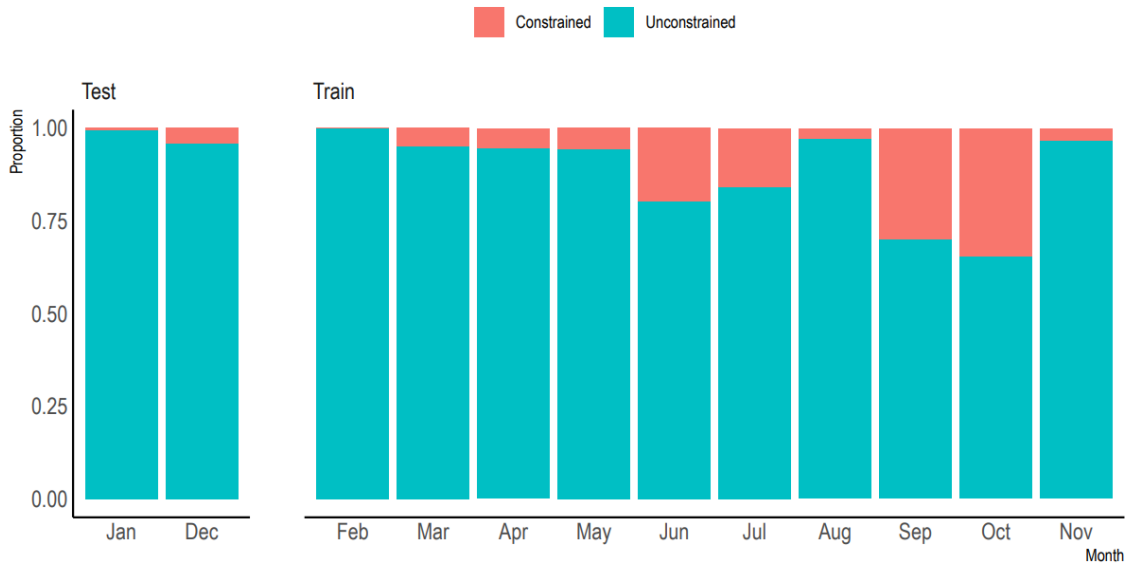

Diffuse irradiance, ambient relative humidity, wind speed and wind direction were missing in the data set from February 2019 to April 2019, so these variables were omitted from the model. Other measures of solar irradiance were available over the complete dataset, so diffuse irradiance was considered less important to include. Similarly, temperature was available over the complete dataset and adequately captures the ITB transformers' efficiency thereby reducing the need for humidity and wind speed. Active power was not captured at all in the data set so gross power was multiplied by an efficiency loss factor to derive net power as a proxy of active power.

Sky Camera imaging for the solar farm was initially intended to be included as an input to the solar farm forecasting model but was eventually dropped due technical difficulties with Sky Camera commissioning and logistical problems associated with COVID-related interstate travel restrictions.

1.3 Machine learning algorithms

For both Waterloo and Ross River, machine learning forecasting algorithms were developed. Various model variants were evaluated. These included state-of-the-art Recurrent Neural Networks (RNN), ETS and ARIMA models. Computation speed was also a critical factor for the models as it is preferable to generate forecasts at high frequency to enable the forecast to be received as close as possible to the AEMO gate closure.

The final model that was chosen for the wind and solar farms displayed high speed, good accuracy, an ability to manage missing values (i.e. during a constraint period), and demonstrated ease of implementation and deployment, with the ability to be quickly fitted to the data set.

The machine learning algorithms also include many engineered features. Modifications of the raw variables (or engineered features) were added into the models, including rolling statistics using an iterative assessment process. Inputs were individually added to the model, then the

model was trained and tested; those that improved forecasting accuracy based on the MAE and RMSE were kept and those that did not were discarded.

The machine learning algorithms for both farms are setup to run as an ensemble to deal with the challenges associated with missing data when there are AEMO imposed constraints on generation output. During normal operation the model runs using 30 minutes of past historical SCADA data and active power data. When constraints are active, past values of active power are not available to produce a forecast, so the algorithm utilises lags of possible power and other SCADA variables. The third scenario occurs when the Farm is coming out of constraint, but a full 30 minutes of unconstrained history is not yet available. In this instance the algorithm utilises a model with a smaller historical window that also considers ramping immediately after a constraint is lifted.

Finally, each model is independently optimised for MAE and RMSE with the algorithm combining the two together.

The combination of these different models running in parallel means that our algorithm can respond more accurately when constraints are active and lifted by switching between models.

1.4 Forecasting horizon

Whilst the AEMO self-forecast is notionally described as a 5-minute ahead forecast, it should more accurately be described as a 6-minute and 10-second ahead forecast, as forecasts must be submitted 70 seconds before the 5-minute gate closure time.

The forecast horizon also needs to consider the delays associated with receiving the data, processing the model and communicating with the AEMO API; approximately 10 seconds. Since minutely data are used, the models were trained to forecast 6 minutes and 7 minutes into the future, then interpolate the forecasts to obtain a 6-minute and 20-second ahead forecast.

1.5 Model deployment

Each of the forecasting models has been designed to run at the edge on a Dell Edge 5100 Gateway. The edge gateways are integrated with the onsite SCADA systems to feed live streaming data over Modbus protocol into the machine learning models. The edge devices provide local buffering of data and by not transmitting large volumes of live data, the bandwidth needs for internet communication are reduced. It also means the models will continue to run if the internet connection is interrupted for any reason, providing a seamless recovery once communications are re-established. Connection to the AEMO API, remote monitoring of model performance as well as patches and updates to model and edge firmware are all enabled over internet connection.

2. Overview of forecast model performance

2.1 Waterloo Wind Farm model performance

In total, 16,123 5-minute dispatch intervals were evaluated over the assessment period. Out of those, 1,442 were constrained and 14,681 were unconstrained. During the assessment period, our algorithm successfully outperformed the benchmark AWEFS (ANEMOS) forecasts.

A summary of the results is shown below. Our algorithm showed an average improvement of 19.8% on MAE and 45.6% on RMSE during unconstrained periods.

Table 2-1 Improvements after applying Advisian's algorithm compared to AWEFS forecasts

		SF (MW)	AWEDS (MW)	Improvement (%)
MAE	Unconstrained periods	2.498	3.115	19.8
	Constrained periods	3.380	9.180	63.2
	All periods	2.577	3.658	29.5
RMSE	Unconstrained periods	4.026	7.404	45.6
	Constrained periods	5.551	21.870	74.6
	All periods	4.185	9.628	56.5

Figure 2-1 Improvements after applying Advisian's algorithm – Unconstrained Period, MAE

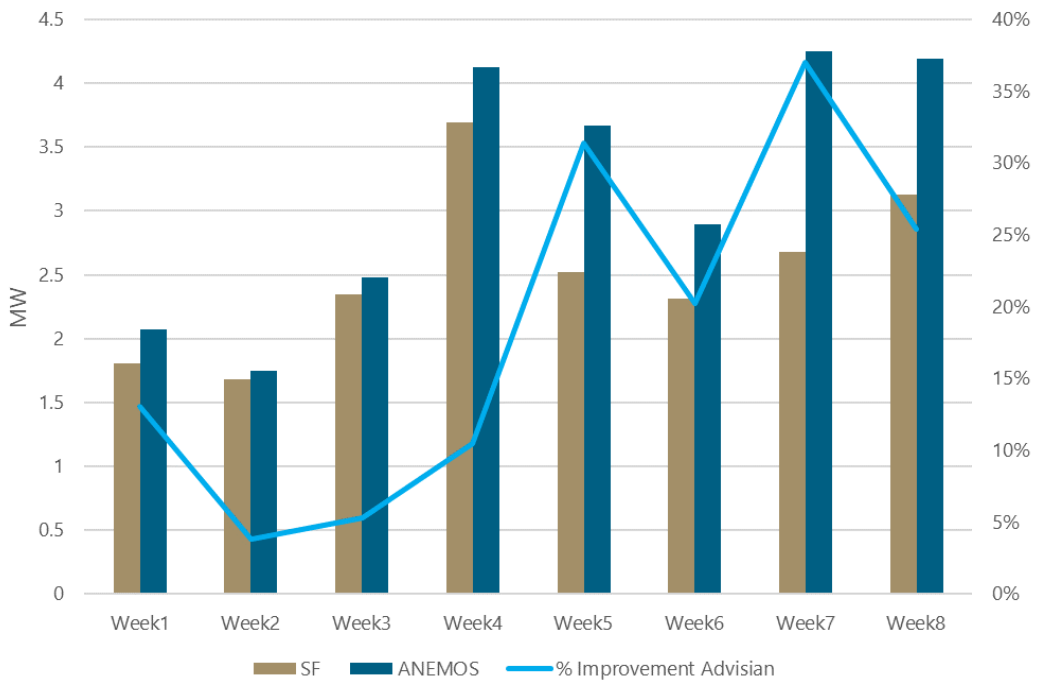

Figure 2-2 Improvements after applying Advisian's algorithm – Unconstrained Period, RMSE

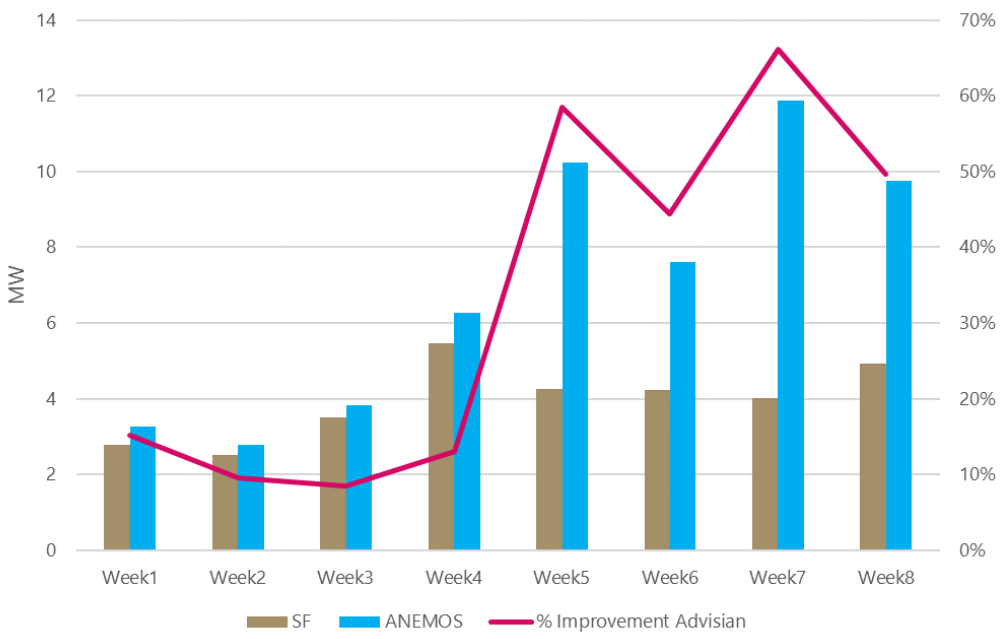

Table 2-2 below shows the measures of Kurtosis and Skewness of the residuals (error) distribution. When compared to the benchmark ANEMOS forecast our model shows significantly lower Kurtosis, and is neutrally skewed whereas the ANEMOS model is heavily negatively skewed.

Table 2-2 Comparison of Kurtosis and Skewness for Unconstrained Periods

	Kurtosis	Skewness
AWEFS	114.28	-8.24
Self-forecast	10.52	-0.01

High Kurtosis is a measure of the weight of the tails in the error distribution. The high Kurtosis for the ANEMOS forecast means that its more likely to deliver high forecasting errors than our self-forecasting model. This can be visualised quite clearly in Figure 2-3 below.

The Skewness of the distribution shows the bias of the distribution towards over or under prediction. Our model is neutrally biased with a Skewness very close to zero. The ANEMOS forecast is heavily negatively biased, meaning that it is more likely to over predict the generation output than underpredict.

Figure 2-3 Unconstrained Period Forecasting Residuals by Week

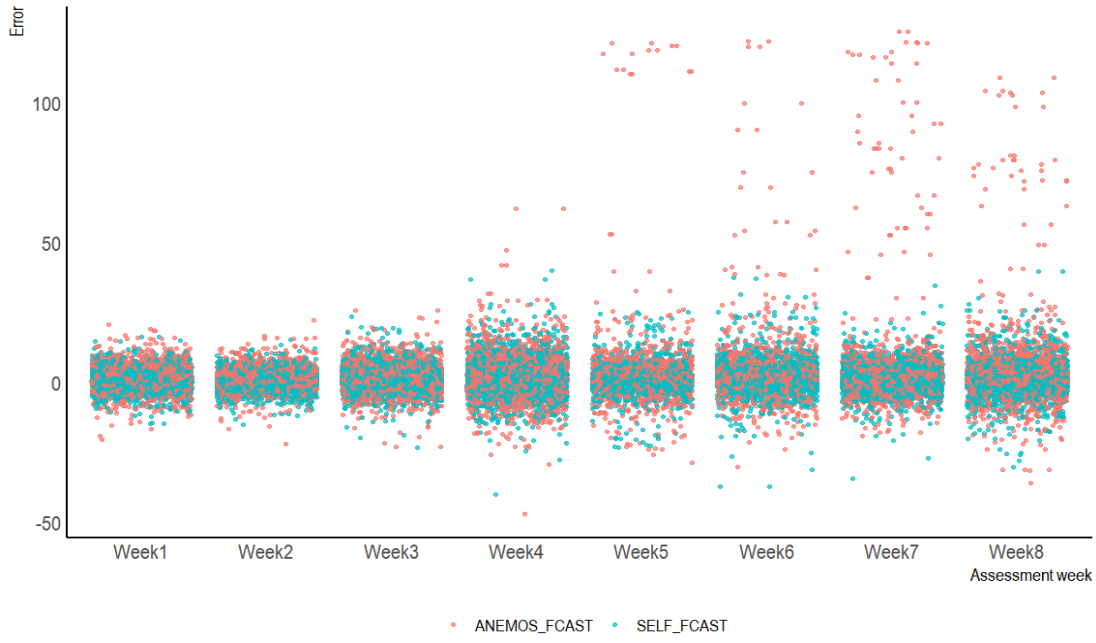

2.2 Ross River Solar Farm model performance

In total, 10,423 5-minute dispatch intervals were evaluated over the assessment period. Out of those, 259 were constrained and 10,164 were unconstrained. During the assessment period, our algorithm successfully outperformed the benchmark ASEFS (ANEMOS) forecasts.

A summary of the results is shown below. Our algorithm showed an average improvement of 2% on MAE and 8% on RMSE during unconstrained periods.

Table 2-3 RRSF Self Forecasting Performance Comparison

		ANEMOS (MW)	SF (MW)	Improvement (%)
MAE	Unconstrained periods	6.4	6.3	2
	Constrained periods	13.0	10.8	17
	All periods	6.6	6.4	3
RMSE	Unconstrained periods	12.9	11.9	8
	Constrained periods	19.3	15.6	19
	All periods	13.1	12.0	9

Table 2-4 below shows the measures of Kurtosis and Skewness of the residuals (error) distribution. Both models perform similarly, displaying a neutral bias and with a relatively low probability of high forecasting errors.

Table 2-4 RRSF Comparison of Kurtosis and Skewness for Unconstrained Periods

	Kurtosis	Skewness
ANEMOS	9.76	-0.19
Self-forecast	8.90	-0.01

Figure 2-4 RRSF Unconstrained Period MAE

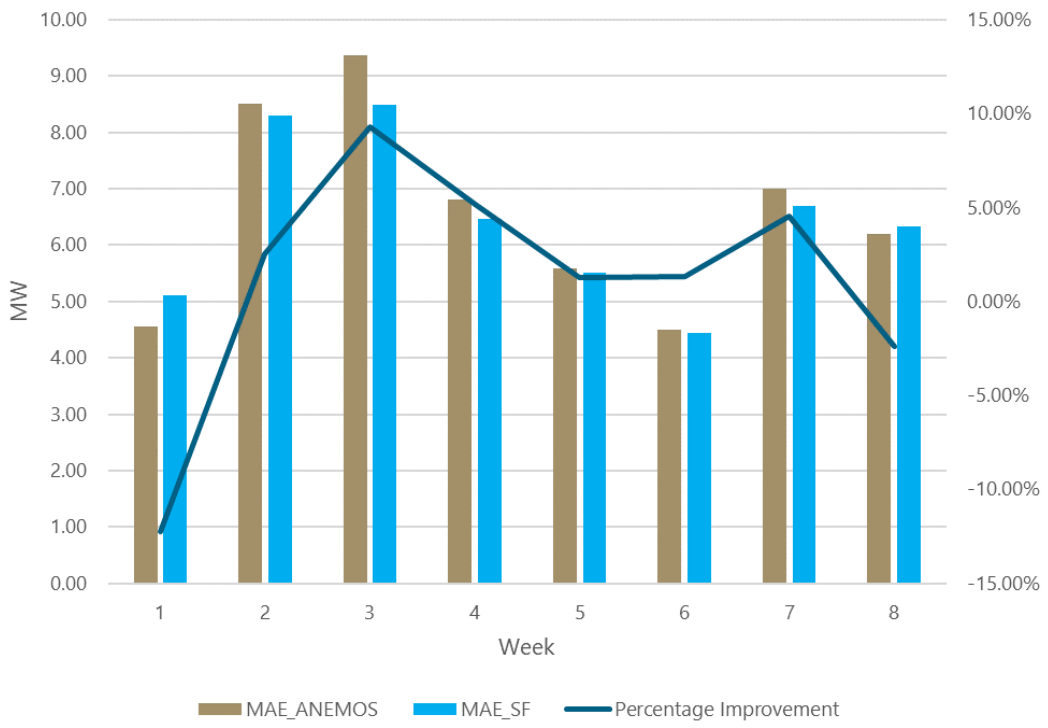

Figure 2-5 RRSF Unconstrained Period RMSE

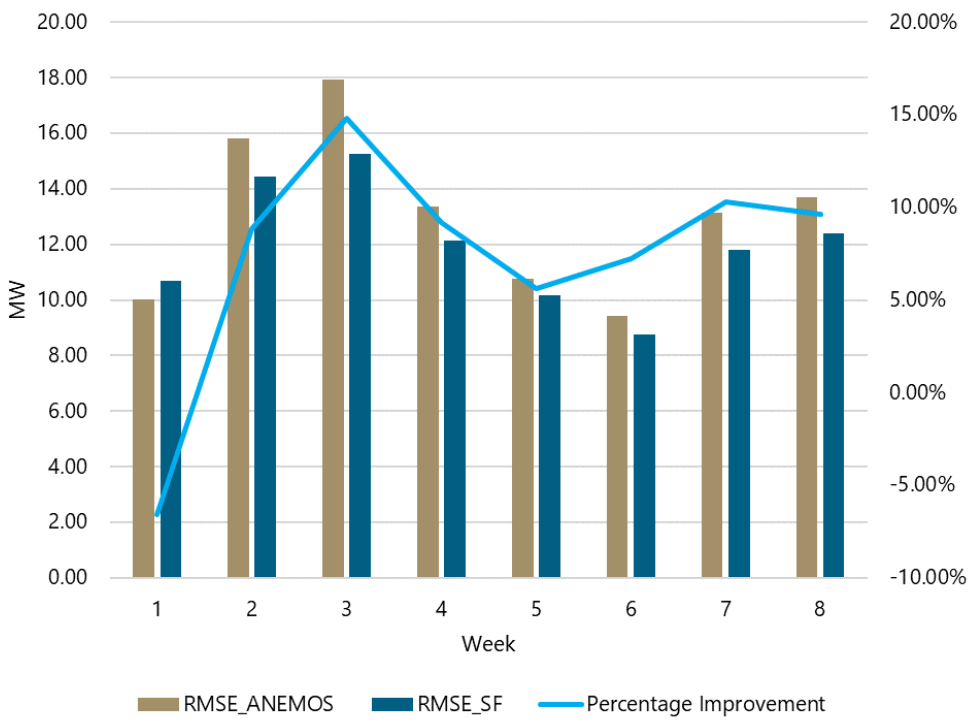

3. Overview of financial benefits

As both RRSF and WWF have existing self-service forecasting providers and do not use the ANEMOS forecast algorithms. It was not possible to benchmark the commercial benefits of our self-forecasting algorithm versus ANEMOS at these sites. To estimate the commercial benefits, historical regulation FCAS (Frequency Control Ancillary Services) charges are divided by total generation capacity to determine an approximate cost per MW/year. Regulation FCAS has historically exhibited extreme volatility as seen in Figure 3-1 below. A conservative approximation of $2,200 per MW/year was used equivalent to $220,000 per year for a 100MW facility. This is aligned with other published estimates. A range of potential savings was calculated based on proportioning the MAE and RMSE improvements to the total annual FCAS cost per MW.

Figure 3-1 Global Regulation FCAS costs by Quarter

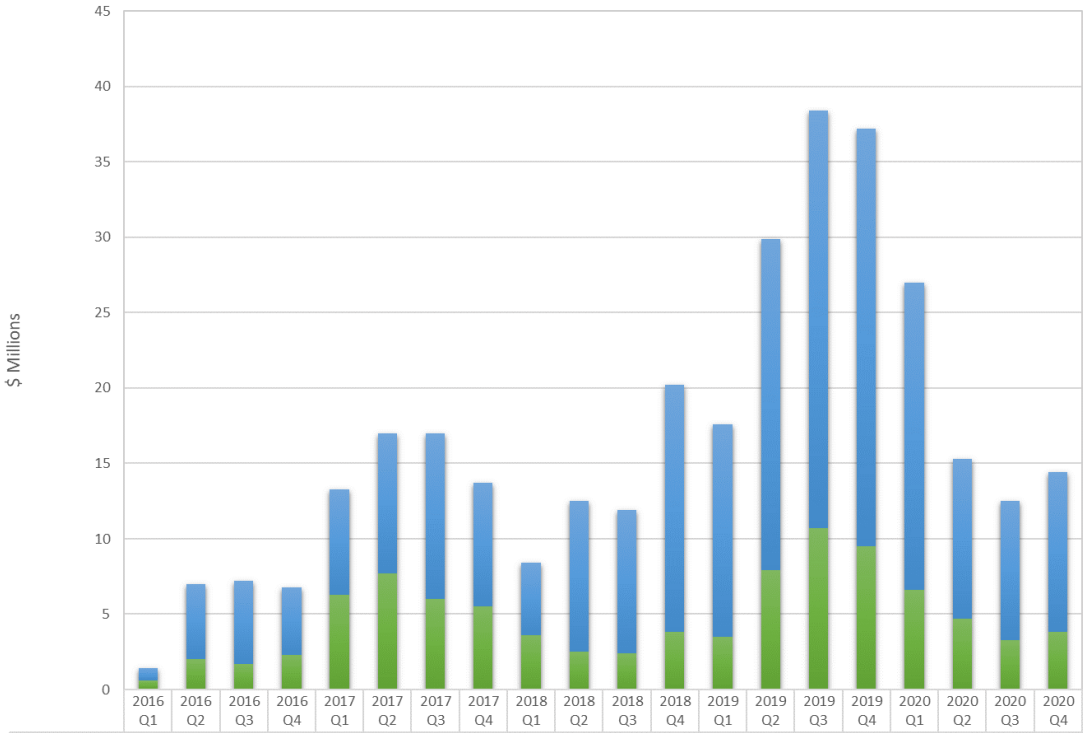

	2016 Q1	2016 Q2	2016 Q3	2016 Q4	2017 Q1	2017 Q2	2017 Q3	2017 Q4	2018 Q1	2018 Q2	2018 Q3	2018 Q4	2019 Q1	2019 Q2	2019 Q3	2019 Q4	2020 Q1	2020 Q2	2020 Q3	2020 Q4
RAISEREG ($)	0.8	5	5.5	4.5	7	9.3	11	8.2	4.8	10	9.5	16.4	14.1	22	27.7	27.7	20.4	10.6	9.2	10.6
LOWERREG ($)	0.6	2	1.7	2.3	6.3	7.7	6	5.5	3.6	2.5	2.4	3.8	3.5	7.9	10.7	9.5	6.6	4.7	3.3	3.8

Based on this approach, annualised savings are estimated between $45,000 to $100,000 per year for the wind power forecasting algorithm, and $5,000 to $20,000 for the solar power forecasting algorithm under 'normal' FCAS price ranges. Savings could easily exceed these estimates, particularly during extreme pricing events such as the South Australia separation incident in late January 2020 where even minor improvements in CPF would yield disproportionately large savings.

One of the features of our wind forecasting algorithm is that it is far less prone to over prediction (Kurtosis and Skewness) than the ANEMOS forecasts. This is particularly evident in the period immediately after constraints are lifted as can be seen in Figure 3-2 below. The benefits of this behaviour are significant but difficult to estimate. Because the Causer Pays Factors (CPF) calculation sums the deviations over the dispatch period, extreme outliers will inflate the net deviation calculation and therefore attract a higher CPF. As a result, the actual benefits to wind farms may be significantly under-estimated.

Figure 3-2 WWF: Difference in errors between ANEMOS and Self Forecast in first period after lifting of a constraint

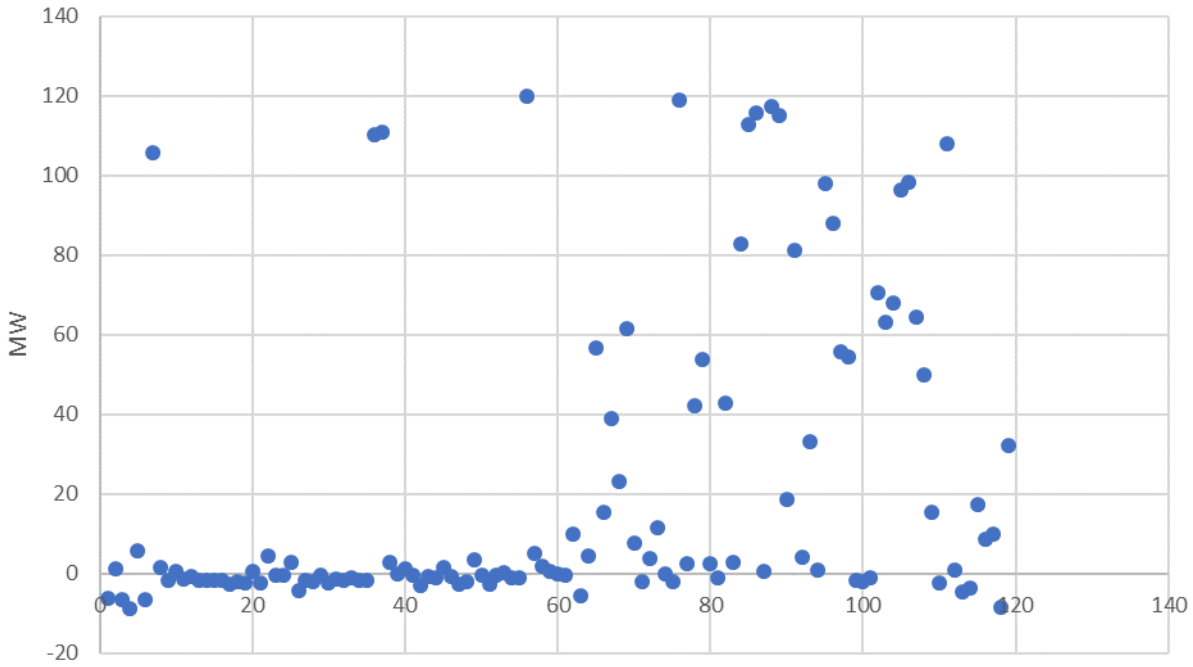

4. Benefits to grid, market and consumers

The objective of regulation frequency control is to achieve generation/demand balance in response to minor deviations in load or generation. Imbalance results in frequency rises or falls which, if not addressed, can compromise overall stability of the system. The FCAS market is the mechanism that AEMO uses to alter the generation and load balance to maintain frequency stability.

The costs associated with providing regulation ancillary services are pooled and distributed on a causer pays basis amongst market participants. Under this methodology the response of measured generators and loads to frequency deviations is monitored and used to determine the CPF.

Poor performance against dispatch targets will generally lead to negative CPF and therefore an increase in a participant's share of regulation FCAS. Wind and solar generators have traditionally paid a disproportionate share of the market FCAS regulation costs[1].

To calculate CPF, deviations from a forecast generation straight line trajectory are calculated at 4-second intervals and multiplied by a Frequency Indicator to generate 4-second performance measures. These are then summed over a 5-minute interval, then grouped into categories, with each category aggregated over the 28-day application period. It is not possible to accurately predict the regulation FCAS costs for a single generator due to the variability in both demand for FCAS and its pricing. Nonetheless, regulation FCAS is a significant operating cost for semi-scheduled generators, a factor that must be considered when assessing the feasibility of new renewable energy projects.

With the introduction of self-forecasting, semi-scheduled generators can set their own 5-minute dispatch targets. The forecasting algorithms developed for this project have demonstrated an ability to out-perform the AEMO AWEF and ASEF forecasting models which would result in a lower CPF for individual generators. If these improvements were extrapolated across all solar and wind farms in the NEM, the cost of regulating FCAS services across the market as whole would reduce, making renewable energy projects more attractive to investors and driving down the cost of energy for consumers.

[1] HARLEY MACKENZIE, NEM FCAS causer pays factor issues for wind and solar farms

5. Technology development

The primary technology that has been developed for this project is the forecasting model and associated technology infrastructure. Following the successful completion of the Wind Forecasting solution, it is the intention of Advisian to commercialise the platform and offer it to our customers in Australia.

6. Key lessons learnt

6.1 Additional sensor measurement

The work undertaken to develop the wind power forecasting model has shown conclusively that it is possible to achieve a high accuracy forecast using only standard SCADA data as inputs. Strong results were achieved without adding new sensors such as meteorological stations, wind masts or lidar.

The opposite is true for solar farms, where models built using only met station and SCADA data did not deliver the targeted improvements in forecasting performance. The addition of sky camera cloud imaging data would have added further information to the model which could have improved the accuracy of the forecast.

6.2 Sky cameras

The project procured scientific sky cameras for capturing cloud imagery. Numerous technical challenges were encountered during commissioning, including integrating the camera into the customer's IT network, adjusting the camera firmware for daytime imaging and overheating of the camera control unit. In hindsight, a commercially available digital sky camera specifically designed for daytime imaging in Australian conditions would have provided a simpler and more reliable solution.

6.3 Edge deployment of the models

Deployment of machine learning forecasting algorithms to the edge is an effective means of deployment with no identifiable adverse effects on latency or model performance.

6.4 Ensemble models

Producing a forecast that is accurate across all operating modes of a wind or solar farm is best achieved by running multiple models in parallel, whereby the algorithm switches between models based on available input data.

6.5 Historical data for self-forecasting

Each forecasting model has been trained to reflect the configuration and operating conditions associated with that site. The forecasting algorithm is designed to be easily replicated to other sites but does require the machine learning components to be retrained on the site's historical data. Many sites historise data infrequently (sometimes 5 or 10-minute intervals) which would leave insufficient information to deliver a forecast to required accuracy. Operators intending to move to self-forecasting should immediately begin historising their data at least 1-minute intervals.

7. Future work

The Raise component of regulation FCAS charges typically makes up the majority (55%-80%) of global regulation FCAS. This may be partly due to the bias observed in the AWEFS model which had a high probability of over predicting generation output, requiring regulation raise FCAS to meet the shortfall. This observation raises an important question; is the project objective of achieving an accurate forecast aligned with achieving lower FCAS costs for generators?

To answer this question, it is necessary to consider whether the self-forecast is working to assist or degrade frequency stabilisation. The forecasting algorithms developed for this project do not incorporate the frequency trend or its inertia within the models. Adding this information may support a forecast that has a bias towards over or under prediction to ensure the generator responds in a manner that generally assists power system frequency. Whilst this would have the benefit of lowering individual generator CPF, biasing the model goes against the objectives of this project which is to develop an accurate 5-minute ahead forecast.

8. Conclusions

The results from this project have demonstrated that improvements can be made to forecasting accuracy for both wind and solar generators by utilising best practice machine learning techniques.

The results from the wind farm was particularly positive, with significant improvements in both RMSE and MAE. The wind model also displayed a robust ability to maintain accurate predictions even during ramping events, such as after constraints were lifted. This is a key feature of the model which will greatly reduce the causer pays factors for generators and provides a level of insurance against severe FCAS costs during extreme network events where regulation FCAS prices are inflated.

The solar model did not achieve the same accuracy improvements as the wind model. Integration of sky camera imagery into the model would likely have led to an improvement in forecasting accuracy.

The objective of this project was to explore methods for self-forecasting that could develop a more accurate 5-minute ahead forecast. Whilst this objective has been achieved there is a question as to whether this is the right objective. Achieving an accurate forecast may not necessarily support frequency stability in the NEM if the 4-second CPF performance measures are working against frequency support. Further work is needed to investigate how to integrate frequency direction and inertia into the model to develop a forecast which aids power system frequency stability.

For more information please contact Denis Marshment (denis.marshment@worley.com)