Market Performance Quarterly Review

October-December 2021

Contents

Figures

1 Purpose of this report

1.1 This document covers a broad range of topics in the electricity market. It is published quarterly to provide visibility of the regular monitoring undertaken by the Electricity Authority (Authority).

1.2 This report also includes:

(a) Using Machine Learning Techniques to Forecast National Electricity Demand.

2 Highlights

2.1 Weekly load in the December quarter was slightly less than the historic average[1], possibly due to partial lockdowns in parts of the country. Weekly load also decreased over the quarter, due to warmer summer weather and the end of year public holiday period.

2.2 Market share of large retailers has increased while market share of small-medium sized retailers has decreased, with Contact showing the greatest growth. Move in switching numbers have grown while trader switches have dropped. These trends are likely due to Covid restrictions loosening for certain parts of the country after being in multi-month lockdowns from August 2021. Trader switches on the other hand declined, possibly as consumers were otherwise occupied with the upcoming summer holiday period.

2.3 Wholesale electricity prices remained relatively low when compared to the rest of the year, due to a combination of increased renewable generation and reduced seasonal demand. Spot prices followed the pattern of decreasing when grid demand decreased, wind generation increased, and gas and coal powered generation decreased.

2.4 Wind generation was unusually high for the quarter, exceeding thermal generation.

2.5 Despite below average inflows, total controlled hydro storage increased to its highest point of the year with most major lakes above their historical mean storage by the end of the quarter.

2.6 Gas demand peaked, adding marginal gas into the market for the first time in months and lowering gas prices.

2.7 Overall forward prices saw minimal change, however, forward prices in mid-2022 were priced at around ~$150/MWh, indicating that the market was factoring in higher than usual supply risk for winter 2022.

2.8 Carbon prices in December's ETS auction reached $68.20/tonne with ensuing secondary market prices exceeding the 2022 auction price cap of $70/tonne.

3 Demand

3.1 Covid restrictions continued to apply this quarter. October saw Alert Level 3 in place around the upper-central North Island region while the remainder of the country was at Alert level 2 'Delta'. On 11.59pm 2 December 2021 the Alert Level system was replaced by the COVID-19 Protection Framework (the traffic light system based on the vaccine pass). Initially Auckland, Northland and a large portion of Whanganui and Rangitikei to East Cape were set at red while the remainder of the country was set to orange.

3.2 Figure 1 shows total daily national grid load over the 2021 December quarter against average daily national grid load for the 2016 to 2020 December quarters. Annotations display the weekly percentage difference between the 2021 load and the 2016-2020

[1] Averaged over October to December for 2016-2020

historical average load.

3.3 Weekly differences in 2021 load compared to historic mean ranged between -4.2 per centand 0.1 per cent. The largest difference of -4.2 per cent occurred in the last week of December.

3.4 Weekly load decreased from 675 GWh at the beginning of the quarter to 586 GWh by the end of the quarter. The trend of decreasing load can be attributed to a shift in seasons as the country transitioned into summer, with less heating and lighting required due to warmer weather and longer daylight hours.

3.5 Reconciled demand for each month October, November and December 2021 was 3,410 GWh, 3,276 GWh, and 3,243 GWh respectively. Commercial demand was around ~78 per cent greater than residential demand throughout the quarter.

Figure 1: Daily Grid Load, Dec 2021 Quarter vs Historic Avg.

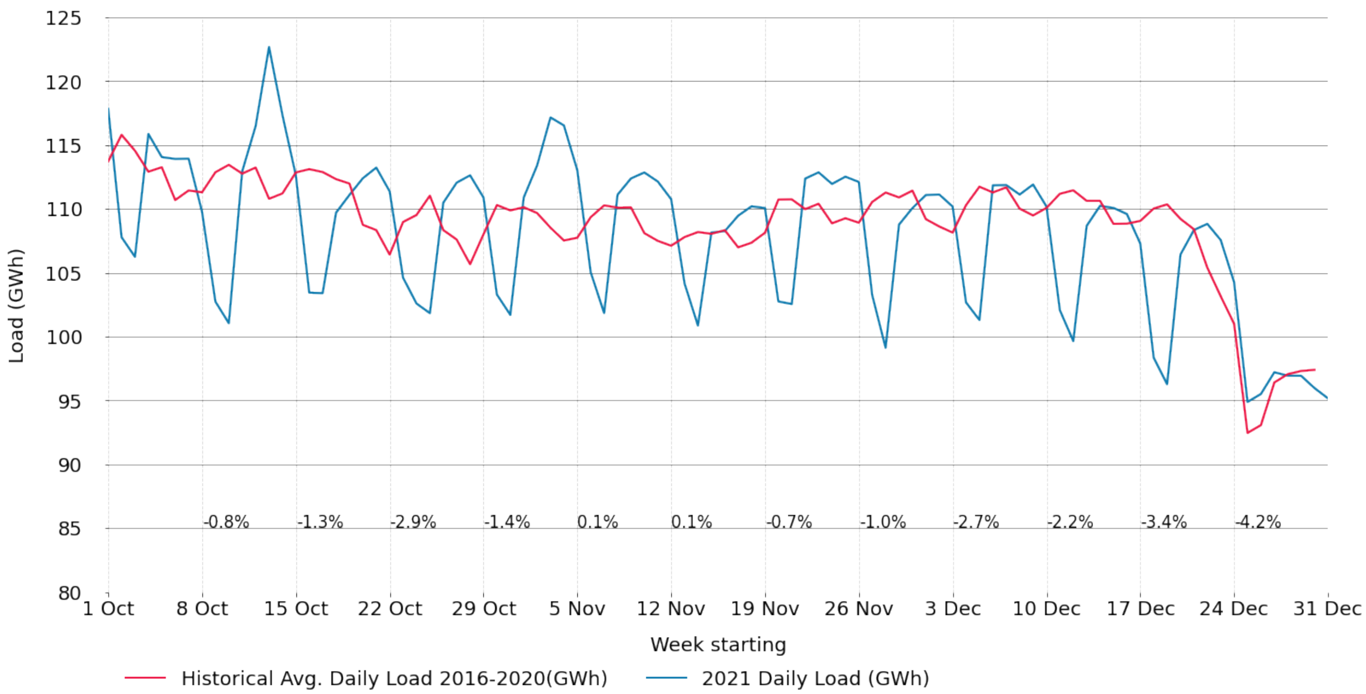

4 Retail

4.1 Market share and participant numbers remained similar to the previous quarter. Over the December 2021 quarter collective market share of the five largest retailers Contact, Genesis, Mercury, Meridian and Trustpower increased minimally by 0.02 per cent to 84.2 per cent.

4.2 On 31 December 2021 the five largest retailers held 1,878,213 ICPs between them, gaining 10,975 ICPs over the quarter while remaining small-medium sized retailers held 352,509 ICPs between them, losing 1,298 ICPs over the quarter.

4.3 Figure 2 shows the changes in market share of each retailer from 1 October 2021 to 31 December 2021. Contact showed the greatest growth gaining 0.23 per cent of market share over the quarter, well above second placed Powershop (parent company Meridian) which gained 0.09 per cent. The retailer with the greatest loss in market share was Genesis with a loss of 0.13 per cent.

4.4 The largest change in ICPs by region by far came from Contact's gains in Auckland with an increase of 4,002 ICPs over the quarter.

Figure 2: Changes in Retailer Market Share[2]

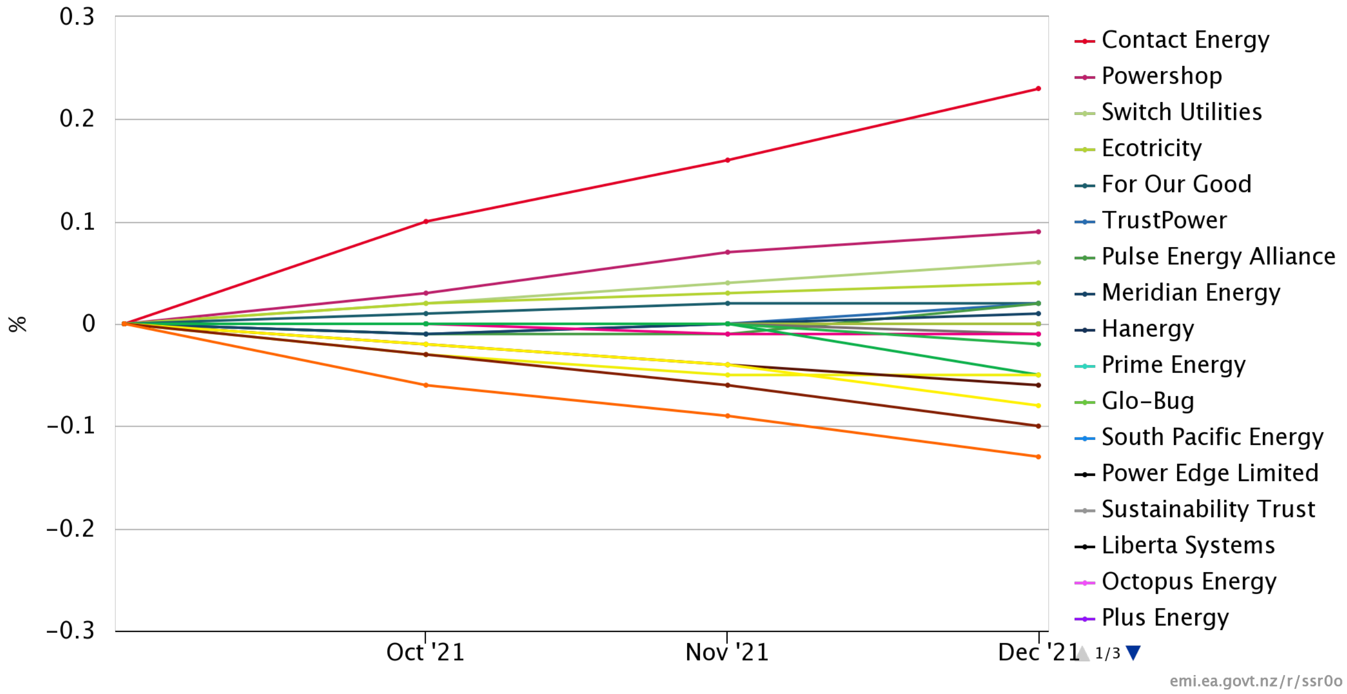

4.5 Figure 3 shows the number of electricity connections (ICPs) that have changed electricity suppliers from 1 October 2021 to 31 December 2021 categorised by type 'move in', 'trader' or 'half hour'. Move in switches are switches where the customer *does not have* an electricity provider contract with a trader, whereas trader switches are switches where the customer *does have* an existing contract with a trader, and the customer obtains a new contract with a different trader.

4.6 Over the quarter move in switches climbed, likely due to Covid restrictions loosening for certain parts of the country after being in multi-month lockdowns from August 2021. Trader switches on the other hand declined, probably as consumers were otherwise occupied with the upcoming summer holiday period. Move in switches increased by 3,805 ICPs while trader switches declined by 1,893 ICPs.

4.7 Compared to the previous five years total trader switches were at their lowest though total move in switches were above average.

[2] Please note that not all traders fit in the key, please go to emi.govt.nz/r/ssr0o to view key with all traders

Figure 3:ICP Switches by Type 2021

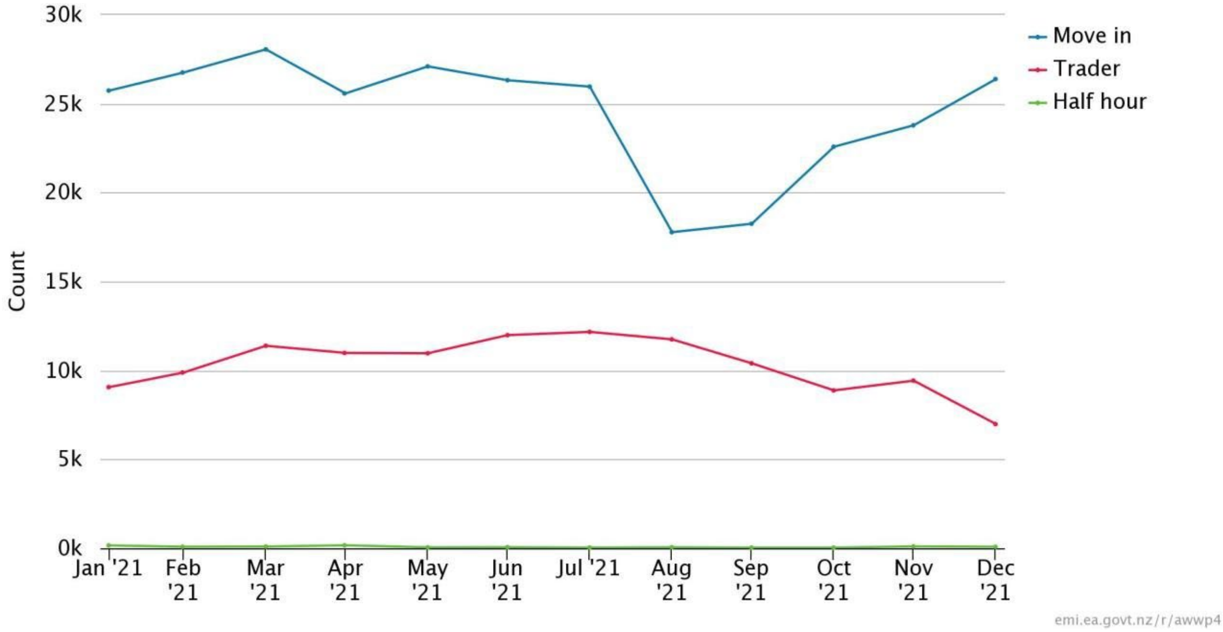

4.8 All three conditions of sale were met for the sale of Trustpower's retail business to Mercury with the High Court's decision in mid-December to approve the proposed restructuring of Tauranga Energy Consumer Trust's (TECT) trust deed. The acquisition is expected to be completed by the first quarter of 2022.

5 Wholesale

5.1 The majority of wholesale electricity spot prices remained below $200/MWh for the December quarter. Spot prices declined from the September quarter due to a higher percentage of renewable generation and reduced seasonal demand.

5.2 Wholesale electricity spot prices from 1 October 2021 to 31 December 2021 are shown below in Figure 4. Spot prices averaged $68.54/MWh across all nodes with 95 per cent of prices falling between $0.01/MWh and $158.90/MWh.

5.3 Following historical patterns, the end of year public holiday period when grid demand is at its lowest averaged the lowest prices for the quarter, around ~$27/MWh from 17 December 2021.

5.4 The highest prices for the quarter of just above $550/MWh occurred due to some combination of above average peak demand, low wind generation, unexpected generation and lines outages and a steep offer curve.

5.5 Outages of note for the quarter include work for the Clutha Upper Waitaki Lines Project (CUWLP) which began on 4 October 2021. Consequent branch constraints at circuits from Clyde to Twizel and Naseby to Roxburgh caused periods of price separation in the region.

5.6 Generally, outside of unusual circumstances spot prices decreased when grid demand decreased, wind generation increased, and gas and coal powered generation decreased. An increase in hydro storage also helped to lower prices by lowering hydro generation costs. Thermal generation offers set high prices for higher tranche offers due to a constrained gas supply keeping the cost of thermal generation high.

Figure 4: Average Daily Wholesale Spot Prices

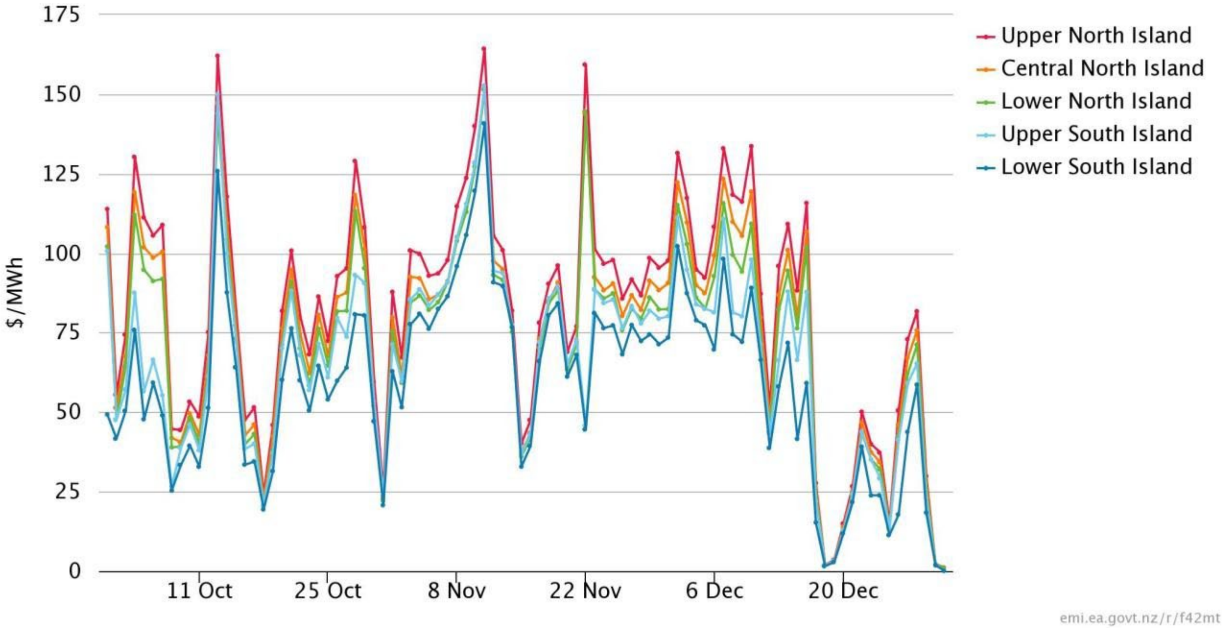

5.7 Generation from renewable sources for the quarter averaged just over ~90 per cent of total generation. The rise in renewable generation has come from a combination of strong wind including extra capacity from Turitea North Wind Farm, steady geothermal and above average hydro.

5.8 Figure 5 shows daily generation for the quarter by fuel type. Wind generation showed strong numbers with half hourly generation averaging 314 MW for the quarter or 7.3 per cent of total generation. Thermal peaker generation conversely was quite low, averaging 16 MW or 0.3 per cent of total generation. Half hourly hydro generation averaged 2,900 MW (65 per cent of total generation) and half hourly thermal generation averaged 258 MW (5.7 per cent of total generation). Unusually, total wind generation exceeded total thermal generation for the quarter.

Figure 5: Average Daily Generation by Fuel Type

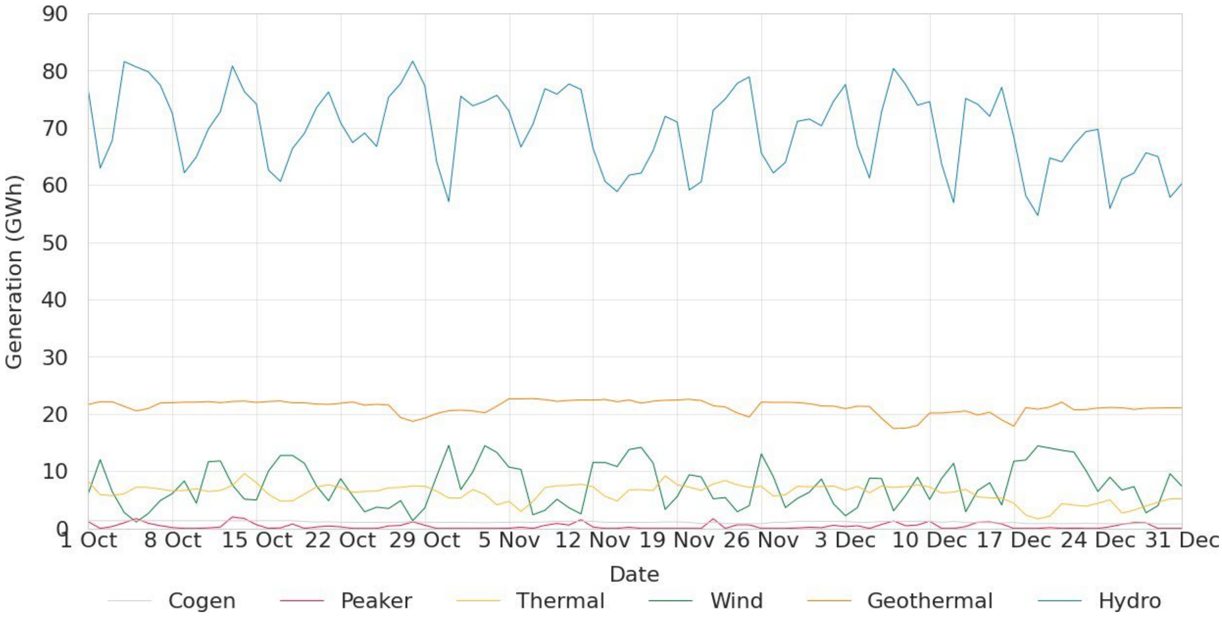

5.9 Hydro generation was aided by an increase in hydro inflows and storage. Figure 6 shows total national controlled hydro storage for 2021. Over the December quarter hydro storage rose by 646 GWh from 3,503 GWh on 1 October 2021 to 4,149 GWh on 31 December 2021. On 31 December 2021 hydro storage was at its highest point for the quarter as well as the year at 125 per cent of historical mean (3,315 GWh) and 94 per cent of nominal full (4,412 GWh).

5.10 Though by historical standards hydro inflows were not especially strong for the quarter compared to earlier in the year inflows showed improvement with hydro storage also aided by reduced demand. Over the December 2021 quarter daily inflows averaged 79 GWh for a total of 7,272 GWh. In comparison historical 1926-2020 December quarter data had daily average inflows of 82 GWh for a total of 7,618 GWh.

Figure 6: Controlled National Hydro Storage 2021

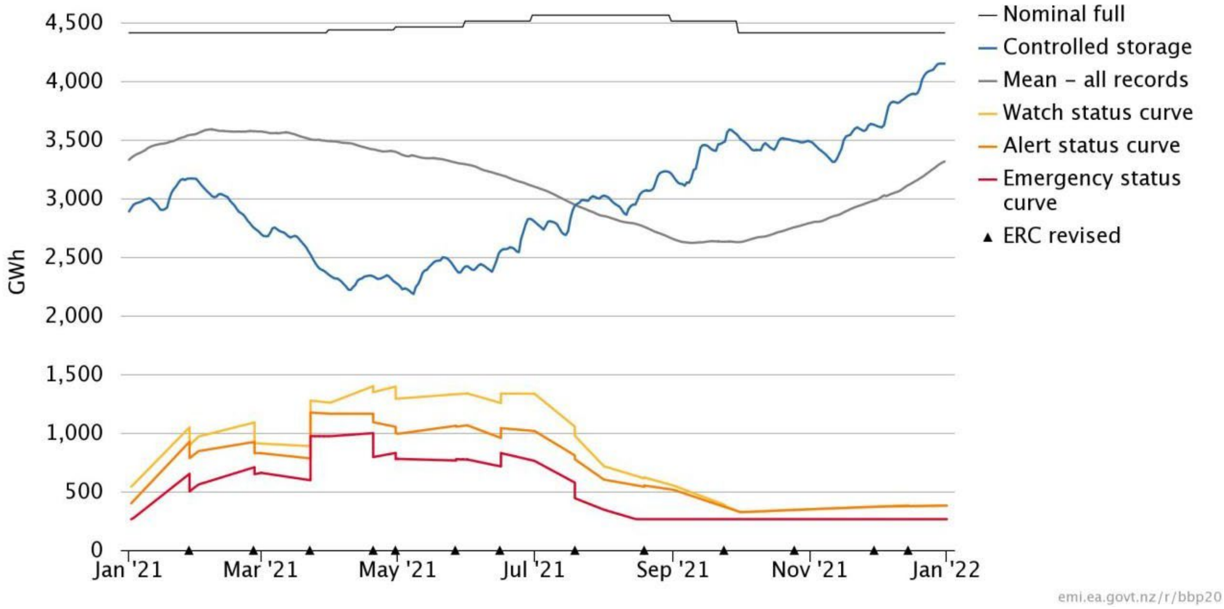

5.11 Figure 7 shows the storage of major catchments Hawea, Manapouri, Pukaki, Taupo, Te Anau and Tekapo for the quarter against their historical means and 10th-90th percentiles based on data from 1926-2022. With the exception of the Manapouri/Te Anau schemes the storage levels of major lakes over the quarter were quite healthy at well above their historical means and increasing close to their 90th storage percentiles.

5.12 In the latter half of the quarter Lake Tekapo was forced to spill water to avoid going over its maximum operating range of 710.9m above sea level. The spill came from a build-up of water in the lake which was a result of Genesis upgrading its Tekapo B station from early October 2021 which halved its generation capacity to 80MW.

Figure 7: Major Lake Storage v mean, 10th-90th percentile

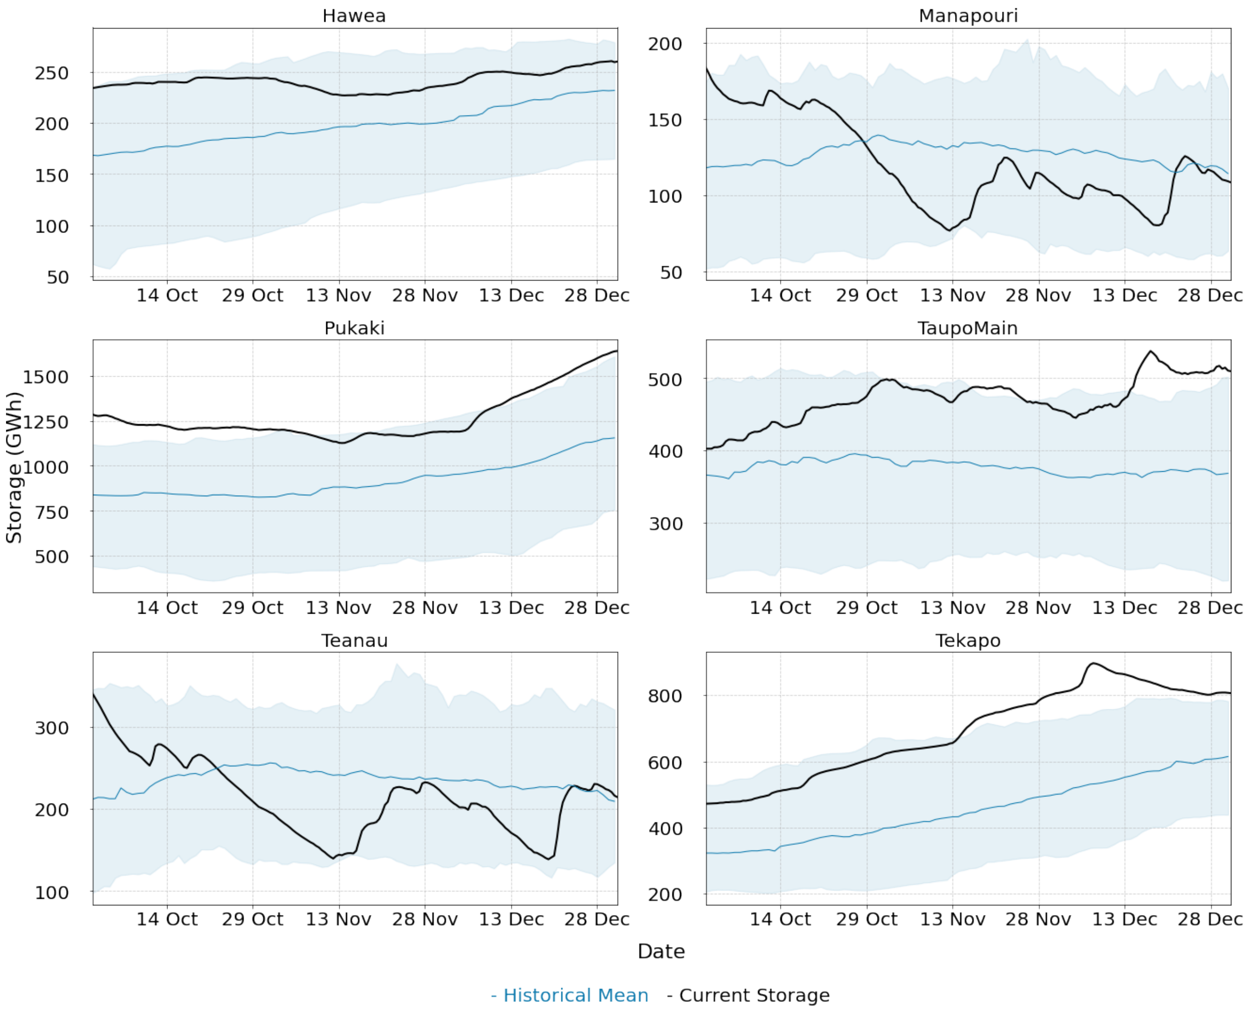

5.13 Figure 8 shows gas production by major fields and gas consumption by major users for 2021.

5.14 Total gas production for the December quarter increased by 12.93 TJ/day from 392.52 TJ/day on 1 October 2021 to 405.43 TJ/day on 17 December 2021 (before gas demand decreased due to the effects of the public holiday period).

5.15 Maui became the largest producing gas field following the successful drilling of the sixth well of OMV's Māui A Crestal Infill drilling campaign, which added over 30 TJ/day to gas production. Gas production at Māui on 17 December 2021 was 100.58 TJ/day.

5.16 Production at Pohokura showed instability for most of October, following an extended outage in September, and fluctuated between ~70 TJ/day and ~100 TJ/day for most of the quarter.

5.17 Consumption by major gas user Methanex increased. Gas consumption was high for most of the quarter peaking at 188.65 TJ/day in late December, up from ~95 TJ/day it was consuming in winter when it let its Motunui plant idle as part of a gas swap deal with Genesis (due to low hydro levels Genesis bought between 3.4PJ and 4.4PJ of gas from Methanex which it used to support thermal generation in winter). In December Methanex's two plants were operating at or near capacity resulting in some surplus gas in the market for the first time since winter 2021.

5.18 Along with Huntly demand steadily declining, due to reduced demand for thermal gas generation gas, suppliers added marginal gas into the market. First Gas and related parties took advantage of this by increasing storage at Ahuroa.

Figure 8: Daily Gas Production and Consumption1 2021[3]

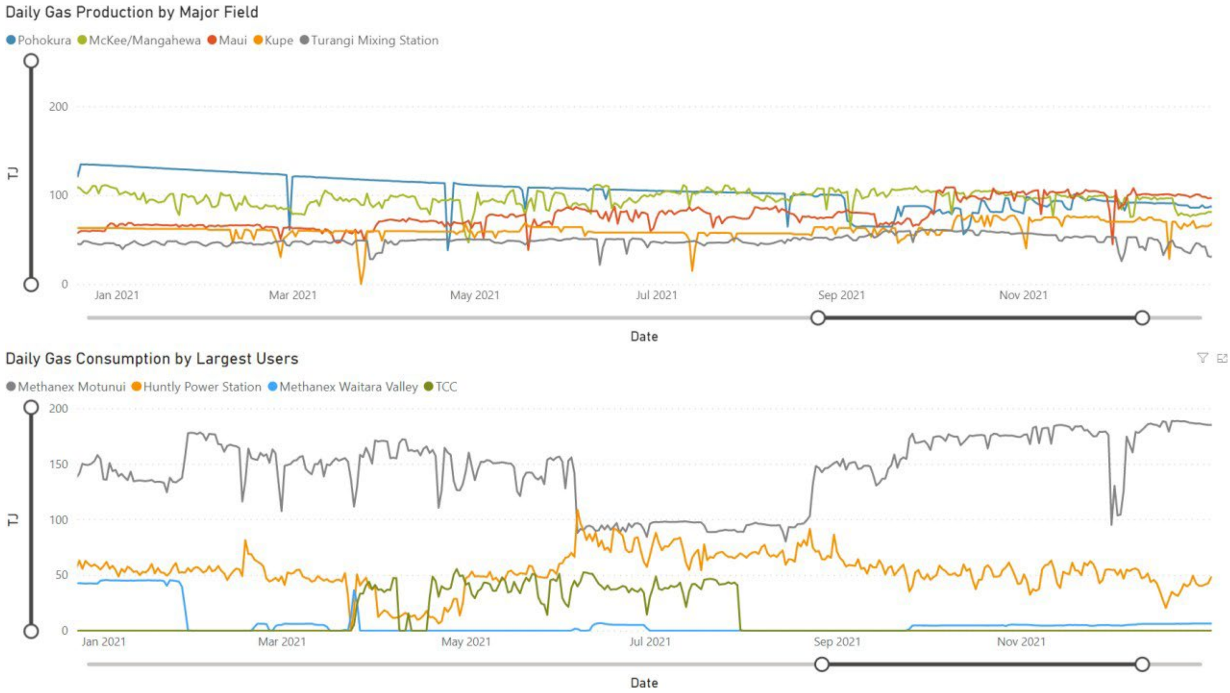

5.19 Figure 9 shows the Maui pipeline average marginal price (AMP) for the December quarter. Pricing data was taken from BGIX2 (Balancing Gas Information Exchange) which we use here as a proxy for gas spot prices.

5.20 Following the increase in gas availability gas spot market trading volumes increased and spot prices dropped. Gas prices for most of the quarter fluctuated at around $10/GJ, half of what prices were in winter which often exceeded $20/GJ. Gas volume weighted average price (VWAP) was $11.36/GJ in October, $11.15/GJ in November and $9.50/GJ in December.

5.21 Coal prices conversely rose significantly, with Indonesian coal (the coal Genesis imports to help power its Huntly Rankines) reaching over US$200/tonne at some points of the quarter, almost four times the price in 2020, due to a mixture of factors including Covid. The increase put the opportunity cost of running the Huntly Rankines at over ~$200/MWh.

5.22 Though Huntly's coal stockpile was high at the end of quarter, Genesis numbers put the stockpile at 835,000 tonnes (1,670 GWh), an increase of 327,000 tonnes from September 2021, thermal generation offers tend to reflect the opportunity cost of using thermal fuel with offers based on current market prices rather than costs incurred.

5.23 Notably in December no coal was required for electricity generation.

[3] https://www.gasindustry.co.nz/about-the-industry/gas-industry-information-portal/gas-production-and-major- consumption-charts/

Figure 9: Daily Spot Gas Prices[4]

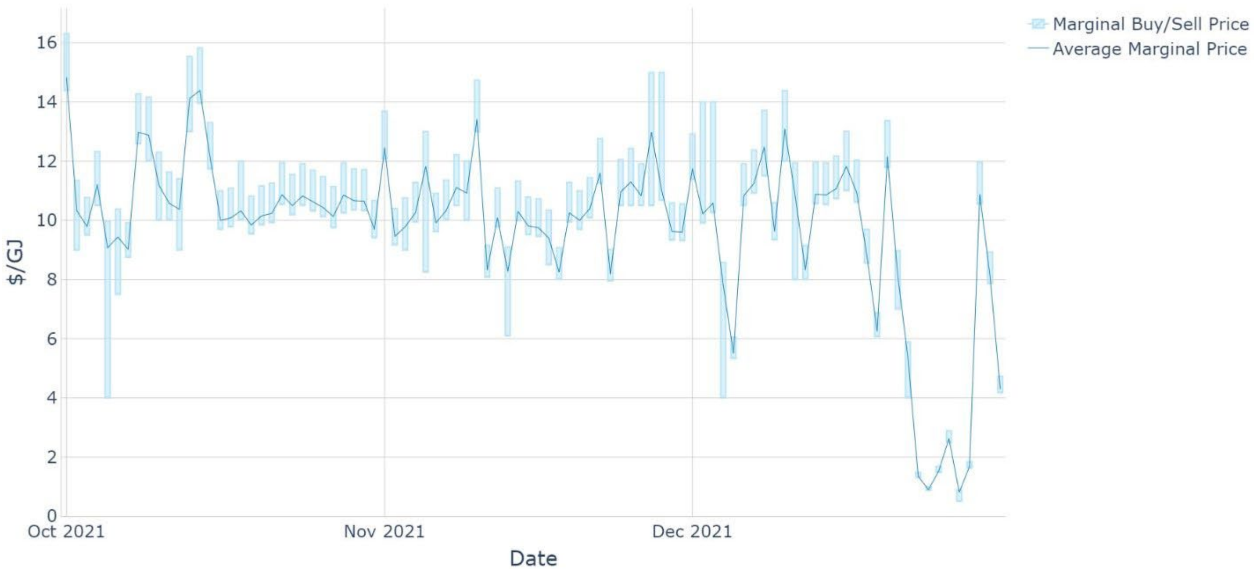

[4] https://bgix.co.nz/prices

6 Forward Market

6.1 The ASX forward price curve provides a view of future wholesale spot prices. Figure 10 shows forward prices for Otahuhu and Benmore at the beginning of the quarter and at the end of the quarter to illustrate how forward prices have changed over the quarter. Forward prices over December saw little change as trading activity took a pause over the public holiday period.

6.2 Short term forward prices (Dec 2021 and March 2022 quarters) rose by around~$18/MWh over the quarter, sitting between $61.10/MWh and $115.25/MWh on 31 December 2021. Long term forward prices either rose by a small amount or remained unchanged, increasing by an average of around ~5$/MWh and sitting between $71.75/MWh and $160.00/MWh on 31 December 2021.

6.3 Forward prices in mid-2022 were priced at around ~$150/MWh showing the market is factoring in higher supply risk for winter 2022. Below average gas production and potential low future inflows at the time increased market participant concerns that the same conditions that caused high prices in 2021 could repeat in 2022.

6.4 The concern for low potential future inflows came from NIWA's November 2021 - January 2022 climate outlook[5] which reported La Niña conditions occurring in the equatorial Pacific during October and moved NIWA to La Niña Alert. (During La Niña rainfall in the lower and western South Island is reduced which usually results in below average inflows and storage levels in South Island catchments.)

6.5 On top of current limited gas production forward prices also factored in the risk outages scheduled in early 2022 at the Māui and Mangahewa gas fields could pose to thermal generation capacity. Maui is scheduled to undergo maintenance for 31 days from 26 March 2022, undergoing a full field outage from 2 April 2022. Mangahewa was scheduled to undergo a 29 day turn around from 1 April 2022, which was later replaced with two new notices for two minor turn downs in March (6 days) and July (8 days). Concerns were that should these outages extend or return at less than full capacity thermal generation availability would be reduced during peak winter demand.

[5] https://niwa.co.nz/climate/seasonal-climate-outlook/seasonal-climate-outlook-november-2021-january-2022

Figure 10: Future Prices

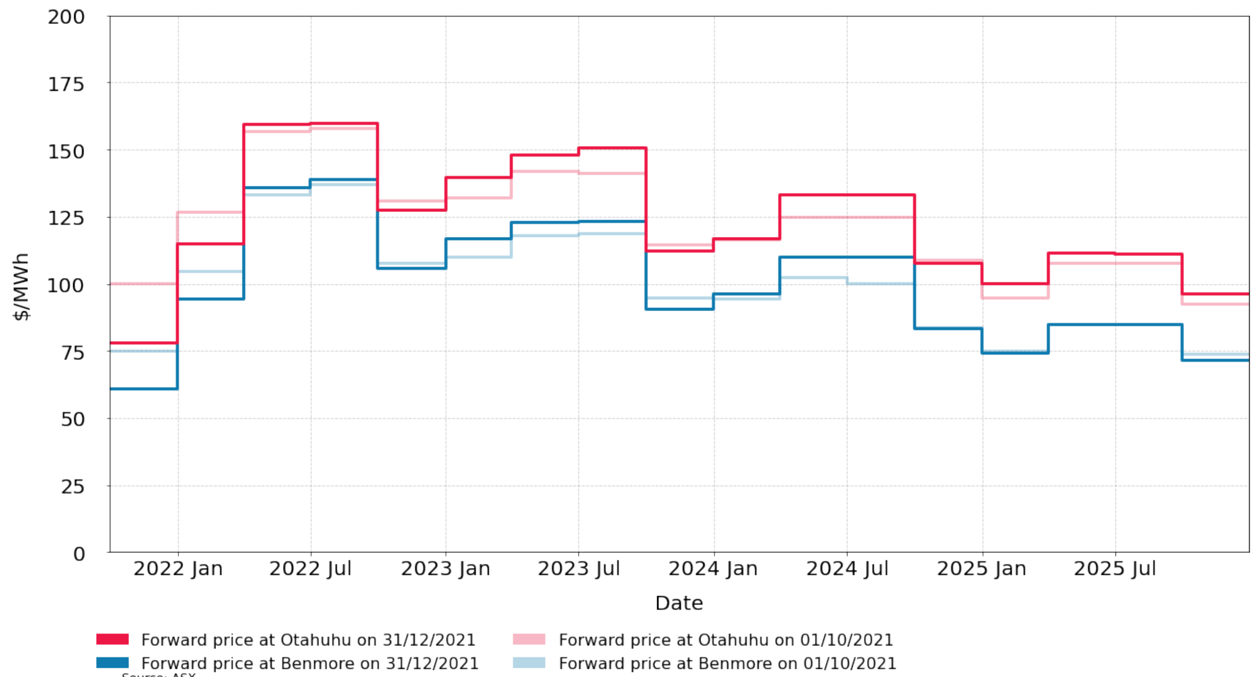

Source: ASX

6.6 Long term forward prices compared to five years ago are noticeably higher. One of the factors that appears to play a role in the rise in forward prices is the increasing price of carbon which increases the costs of running thermal generation.

6.7 Figure 11 shows spot NZ carbon unit prices since the first ETS auction began in 2020, taken from CommTrade's (a platform for buying and selling NZ ETS carbon credits owned by Jarden Securities Limited) website.

6.8 Overall carbon prices rose by 81.3 per cent in 2021. December 2021's ETS auction broke records, rising 5 per cent to $68.20/tonne and ending with 4.75 million units sold with no units left in reserve. Following the auction secondary market prices proceeded to exceed $71/tonne (the auction cap for carbon prices in 2022 is $70/tonne).

6.9 How the price of carbon feeds into thermal generation costs is that approximately ~40 per cent of the price of carbon is added to the final cost of generation for a CCGT (thermal plant) and ~50-60 per cent for OCGTs (peaker plants) when run on gas and ~100 per cent when run on coal. For example, when carbon is at $70/tonne an additional ~$28/MWh to ~$35/MWh would be added to the running cost of thermal generation when run on gas and ~$70/MWh would be added when run on coal.

6.10 The gas outlook for 2023 should be greatly improved compared to 2022, with gas production increases expected at multiple fields. Some of the expected drops in gas prices, however, are likely to be offset by continued increases in carbon prices with the cap for ETS auctions in 2023 set to $78.40/tonne and $87.81/tonne in 2024. It is not clear that generators are facing gas supply constraints next winter, so the increased supply may simply mean more production by Methanex.

6.11 With current renewable generation sources such as wind and solar too intermittent and capacity too low to provide the baseload and firming capacity that thermal generation currently does thermal plants are expected to continue to influence the market for a few more years to come.

Figure 11: NZ Carbon Unit Spot Prices[6]

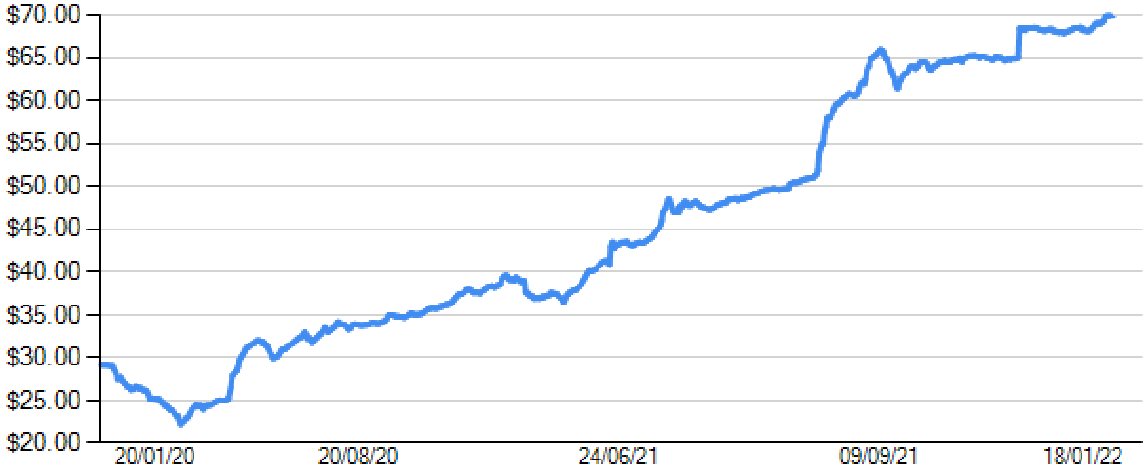

7 Deep Dive: Using Machine Learning Techniques to Forecast National Electricity Demand

7.1 This report includes analysis which uses machine learning to forecast electricity demand in New Zealand, using observational and forecast weather data. The Gluon Time-Series Toolkit (GluonTS), a Python library for deep learning time series modelling, was used forecast national demand. Specifically, a model was built that considers any number of explanatory variables to calculate the national demand in each trading period.

7.2 The analysis found that weather observations in major population centres are the greatest factor that determine demand. To calculate demand one or two days into the future, the best method is to use weather forecasts. This is accurate as long as the forecast and observed temperatures follows the same trends.

7.3 This work on demand forecasting will feed into the response to the Authority's Phase 2 report into the events of 9 August 2021. In particular the Authority is assessing the accuracy of forecasting data that contribute to the pre-dispatch schedules, one of which is the demand forecast.

[6] https://www.commtrade.co.nz/

Using Machine Learning Techniques to Forecast National Electricity Demand

An Overview

Final report

14 December 2021

Market Analytics

Version control

Version	Date amended	Comments
1	14/12/2021	First draft

Summary

We use the Gluon Time-Series Toolkit[1] (GluonTS), a Python library for deep learning time series modelling, to forecast national demand. Specifically, we built a model that considers any number of explanatory variables to calculate the national demand in each trading period. The evening of 29 Jun 2021 was especially cold and national demand peaked at 3.4 GWh. This provides an excellent opportunity to test our machine learning approach. We find that weather observations in major population centres are the greatest factors that determine demand. We explore how to build models in situations where data for some explanatory variables are unavailable. To calculate demand one or two days into the future, the best method is to use weather forecasts. This is highly accurate as long as the forecast and observed temperatures follows the same trends.

1 (i) A. Alexandrov *et al.*, GluonTS: Probabilistic and Neural Time Series Modeling in Python, *J. Mach. Learn. Res.* **21**(116), 2020. (ii) A. Alexandrov *et al.*, arXiv:1906.05264, 2019.

Contents

Tables

Figures

1 Method

1.1 GluonTS is the first ever dedicated toolkit for building time series models based on deep learning and probabilistic modelling techniques. It bundles components, models and tools for time series applications like forecasting and anomaly detection. Well-established open source packages like TensorFlow and PyTorch can perform time-series forecasting, but they require a comparatively large amount of coding and therefore great proficiency in the area of machine learning.

1.2 GluonTS can parse sequences of large collections of time series. Given a probabilistic model, the goal of forecasting is to predict the probability distribution of future values given the past values, covariates (features), and the model's hyperparameters.

1.3 The package provides a wide variety of pre-built neural network based models. For example the DeepAR[2] function uses a recurrent neural network (RNN) with LSTM or GRU cells, and estimates parameters of a parametric distribution or uses a parameterisation of the quantile function.

1.4 The next section is a short overview of the mechanisms behind neural networks[3].

Neural networks in a nutshell

1.5 Artificial neural networks (ANNs) comprise a web of computing units (artificial neurons) organised in homogeneous layers. There is one (passthrough) input layer, one or more hidden layers, and one final output layer. Each layer is connected to the next so that information flows between the layers (Figure 1). Each connection between layers mimics the behaviour of synapses in the brain.

1.6 Abstractly speaking, artificial neurons are considered as functions that take some input values and returns a real number. They have two key roles:

(a) Neurons multiply each input value with a corresponding *weight* coefficient and calculates the sum of all products. This operation is a scalar product of two vectors: input data and weights.

(b) Neurons weigh each input value using an activation function (e.g. logistic sigmoid, tanh) and assigns a quasiprobability value to it.

1.7 The backpropagation algorithm trains the ANN. It is an implementation of the gradient descent method which finds the minimum of a function by exploring values in the direction of the steepest descent. The algorithm makes a prediction at each neuron during the forward pass (going from input layer to output) and measures the error.

1.8 The gradient is calculated on the weights on the output layer, and the error information is pushed backward to the previous layer where the calculation is repeated on the local weights. This recursive process continues until the error reaches the input layer. The aim is to minimise the error of the outputs.

---

2 D. Salinas *et al.*, DeepAR: Probabilistic forecasting with autoregressive recurrent networks, *Int. J. Forecast.* **36**(3):1181-1191, 2020.

3 A. Géron, *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow*, O'Reilly (2nd ed.), 2019.

Figure 1: Simplest model is the feed-forward neural network with backpropagation algorithm. The network can have any number of nodes and layers.

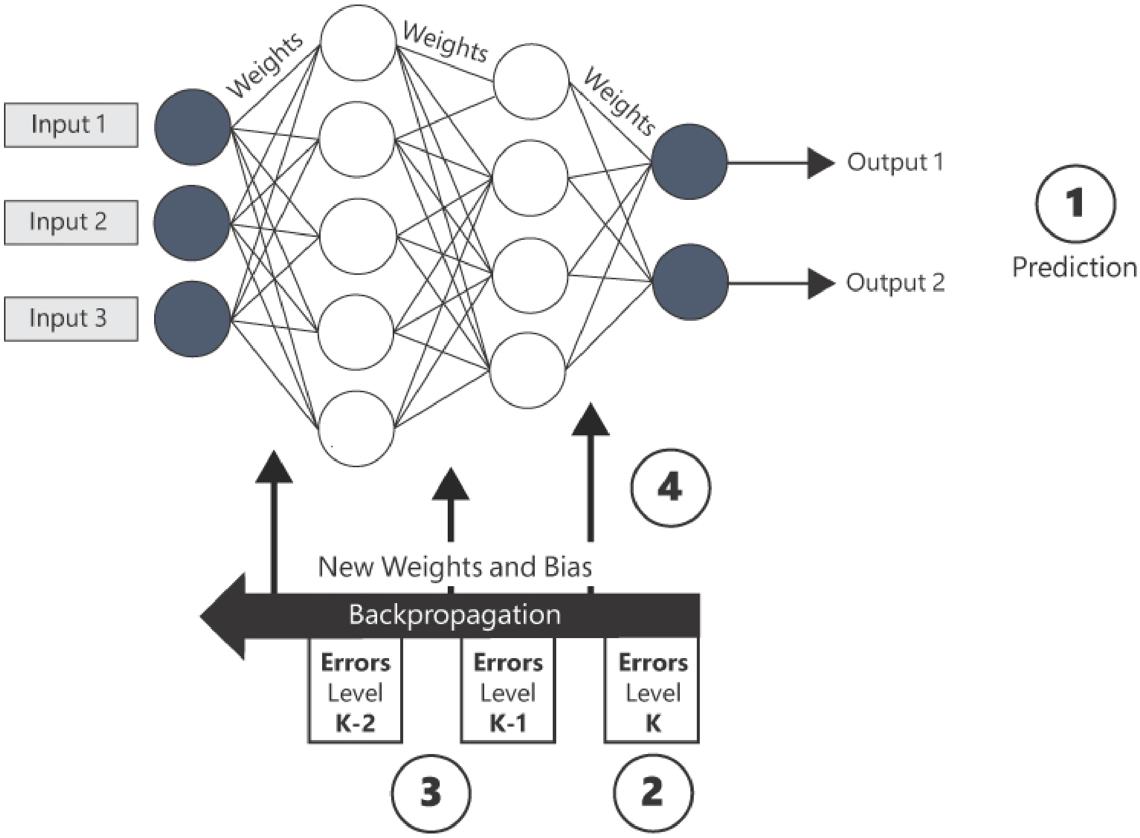

Source: Adapted from D. Esposito and F. Esposito, *Introducing Machine Learning*, Microsoft Press, 2020.

1.9 Time series analysis requires the concept of state to be built inside the ANN itself to make it possible to extrapolate based on historical information. The anatomy of a stateful ANN allows the information to leave tracks as it flows forward from the input layer to the output.

1.10 For this purpose, recurrent neural networks (RNNs) contain hidden states and loops, allowing information to persist over time. Connections between neurons form a directed cycle which creates an internal state and allows the network to exhibit dynamic temporal behaviour. It follows the same concepts as the ANN and uses the same back propagation algorithm.

1.11 RNNs use the idea of hidden state (or internal memory) by updating each neuron so that it remembers what it has seen before (Figure 2). The memory is preserved so that when the neuron reads an input X_n at time step n, it also processes the content of the memory H_n (i.e. outputs of the previous time step $n-1$) and combines the information to generate an output for the current time step Y_n.

Figure 2: Back propagation process in RNNs to compute gradient values.

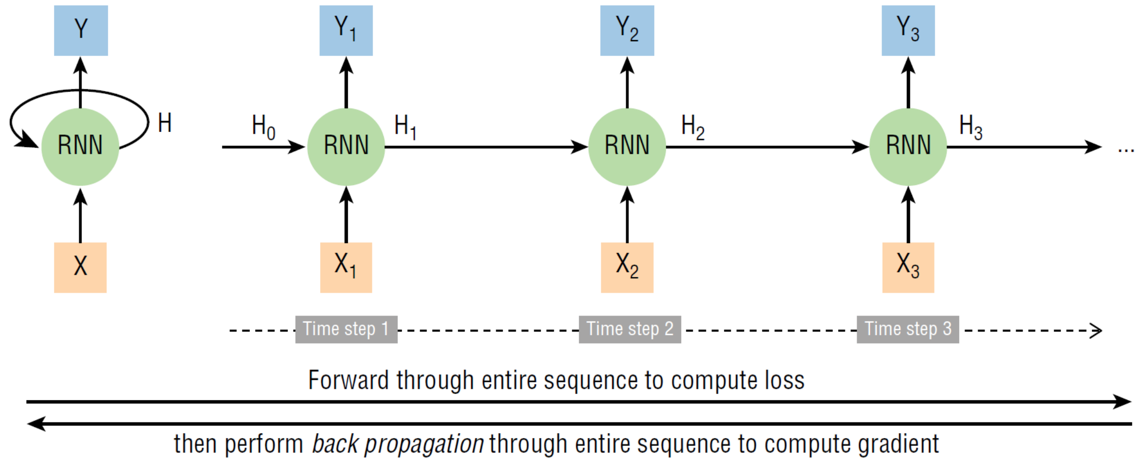

Source: F. Lazzeri, *Machine Learning for Time-Series Forecasting with Python*, Wiley, 2020.

1.12 In order for the cells to remember information from the distant past and not just the previous timestep, the RNN can be built using long short-term memory (LSTM) or gated recurrent unit (GRU) cells (Figure 3). The internal mechanisms in both cells involve learning to recognise an important input (input gate); storing it in the long-term state; preserving it for as long as needed (forget gate); and extract it when needed. Both LSTM and GRU are useful for capturing long-term patterns in time series, long texts, audio recordings, and other applications.

Figure 3: LSTM and GRU cells to replace green RNN cells in Figure 2.

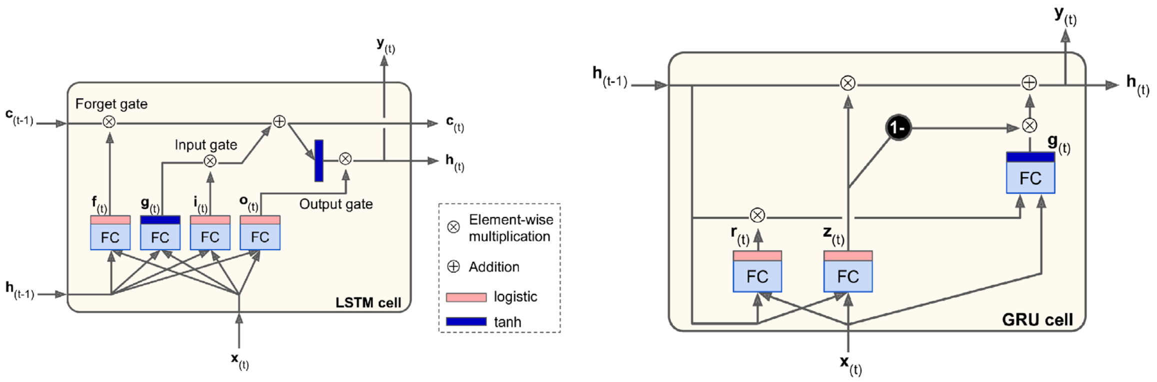

Source: A. Géron, *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow*, O'Reilly (2nd ed.), 2019.

Deep AR model

1.13 In this report, we use a more sophisticated RNN based on Salinas *et al.*'s DeepAR model[4]. It is a forecasting method based on autoregressive RNNs, which learns a global model from historical data of all time series in the dataset (Figure 4).

1.14 In the forecasting community, ANNs are typically applied to individual time series, i.e., a different model is fitted to each time series independently. In comparison, DeepAR tailors a LSTM-based RNN architecture to the probabilistic forecasting problem.

Figure 4: Schematics of DeepAR method used by gluonts.model.deepar.DeepAREstimator function.

Method follows the same structure as in Figure 2. The important addition is p, which is used to calculate the predicted distribution. Further technical details in paper.

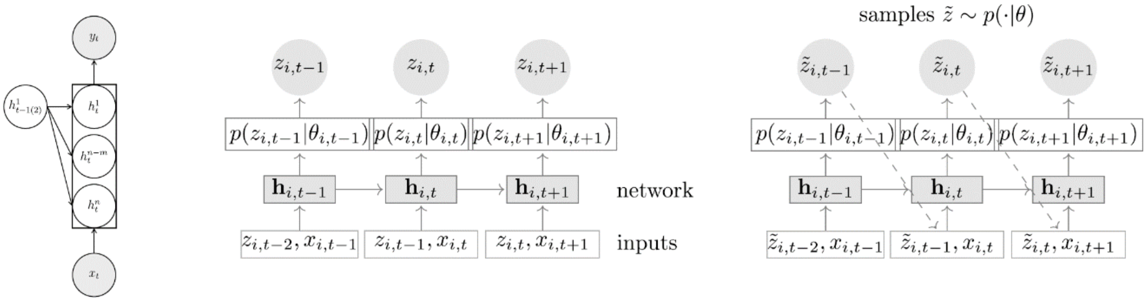

Source: Salinas *et al.*

1.15 As well as having greater forecasting accuracy than classical approaches, DeepAR has several important advantages:

(a) The model learns seasonal behaviours and dependencies on given covariates across time series. Minimal manual intervention in providing covariates is needed in order to capture complex, group dependent behaviour.

(b) DeepAR makes probabilistic forecasts by taking Monte Carlo samples that can be used to compute consistent quantile estimates for all sub-ranges in the prediction horizon.

(c) By learning from similar items, DeepAR can provide forecasts for items that have little or no history available, a case where traditional single item forecasting methods fail.

(d) Our approach does not assume Gaussian noise, but can incorporate a wide range of distribution functions, so we can choose one appropriate for the statistical properties of the data.

1.16 Points (a) and (c) set DeepAR apart from classical forecasting approaches. Points (b) and (d) are important to produce accurate forecast distributions that are learned from

4 D. Salinas *et al.*, DeepAR: Probabilistic forecasting with autoregressive recurrent networks, *Int. J. Forecast.* **36**(3):1181-1191, 2020.

historical behaviour of all time series. This has not been addressed by previous methods.

Data inputs and output

1.17 The idea is to feed inputs into the GluonTS algorithms to see how they affect the national demand.

1.18 Inputs are the apparent temperature[5] in major population centres (Auckland and Wellington, °C). Other inputs were also considered, such as the COVID-19 alert levels taken from government records of key events[6]. However we found that the model had the most accurate results when only given temperature data.

1.19 The output is national demand (GWh)[7].

1.20 We use equally-spaced timesteps of 30 minutes, starting from 1 Mar 2014 00:00:00 which was the first day that half hourly weather data was available (and without any missing or corrupted data). We choose 2 Jul 2021 23:59:59 as the end date, which gives us roughly 130,000 consecutive timesteps (i.e. trading periods) to train and test the model.

Weather data

1.21 Apparent temperature T_A measures what an observer feels[8]. It considers four environmental factors: wind, temperature, humidity, and radiation from the sun:

$$T_A = T_M + 0.348E - 0.70\left(w_s - \frac{Q}{w_s + 10}\right) - 4.25 \qquad (1)$$

with dry bulb (measured) temperature T_M (°C); wind speed w_s (m/s) at an elevation of 10m; net radiation absorbed per unit area of body surface Q (W/m^2); and water vapour pressure (humidity),

$$E = \frac{\mathrm{RH}}{100} 6.105\, e^{17.27\, T_M / \,(237.7 + T_M)} \qquad (2)$$

in hPa where RH is relative humidity (%).

1.22 For simplicity, we assume that the observer is outdoors in the shade so $Q = 0$.

1.23 Instead of considering three environmental factors separately, the apparent temperature is a more compact way to estimate how hot or cold it feels in major population centres (Auckland, Wellington) and therefore the likelihood that many people will switch on air conditioners and heaters.

2 Results

2.1 Here are some examples showing how GluonTS models demand. We build different models that predict demand, based on what data is available.

5 Source: Weather Underground.

6 Source: History of the COVID-19 Alert System, https://covid19.govt.nz/alert-levels-and-updates/history-of-the-covid-19-alert-system/

7 Source: Electricity Authority.

8 R. G. Steadman, A Universal Scale of Apparent Temperature, *J. Clim. Appl. Meteorol.* **23**(12):1674-1687, 1984.

(a) **Many-to-one** forecast models take more than one input and gives one output. This can be used when we know the weather in Auckland and Wellington. It is best suited for checking historical demand since we would have weather and demand data.

(b) **Many-to-many** forecast models take more than one input and gives more than one output. This is used when we do not know the weather and need to predict the demand. We must therefore train the model and ask it to forecast both weather and demand. In other words, this method is used for predicting future trends.

(c) **One-to-one** forecast models take one input and gives one output. Here we would only consider demand.

Many-to-one model

2.2 Figure 5 shows the results of a many-to-one model. Predictions are shown in orange, and actual demand is in brown. The orange line is the median of the model prediction. Probabilistic forecasting requires that we learn the distribution of future values so we need to specify the type of distribution of future values.

2.3 GluonTS comes with many different distributions like Gaussian, Student-t, and Uniform. By default, the model assumes a Student-t distribution and LSTM cells.

Figure 5: Many-to-one model.

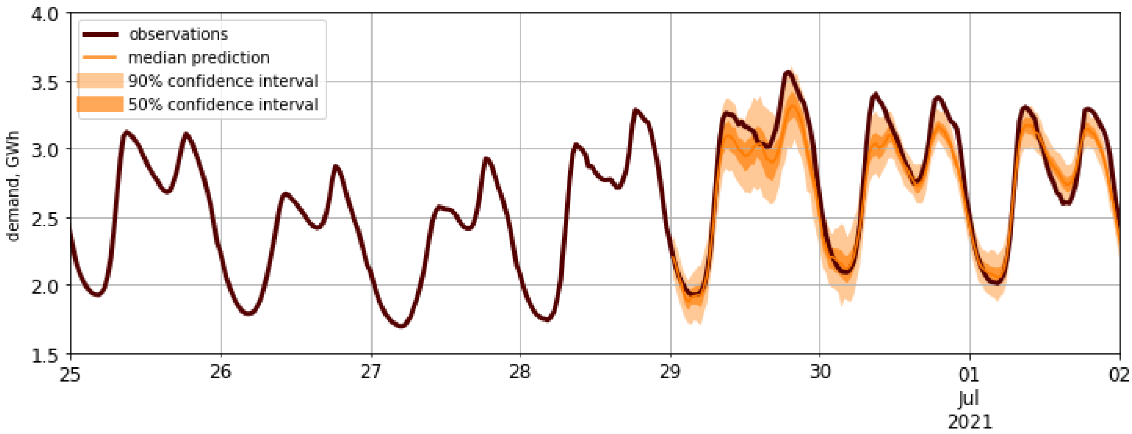

Source: Electricity Authority

2.4 The shaded areas are the confidence intervals surrounding each prediction, in which 90% and 50% of predictions are expected to fall. Smaller intervals imply greater confidence in the prediction.

Many-to-many model

2.5 Figure 6 shows the results of a many-to-many model. Observed values in orange. Model predictions in brown, blue and green. Since we do not know the weather inputs and have to predict demand, the model must calculate forecasts for all three variables.

2.6 As expected, the predictions are poorer than for the many-to-one model. It is caused by the inaccurate temperature predictions for Auckland and Wellington. Because the forecast temperatures are higher than actual temperatures, demand is expected to be lower. This will be a problem if we use this method to predict future demand.

Figure 6: Many-to-many model.

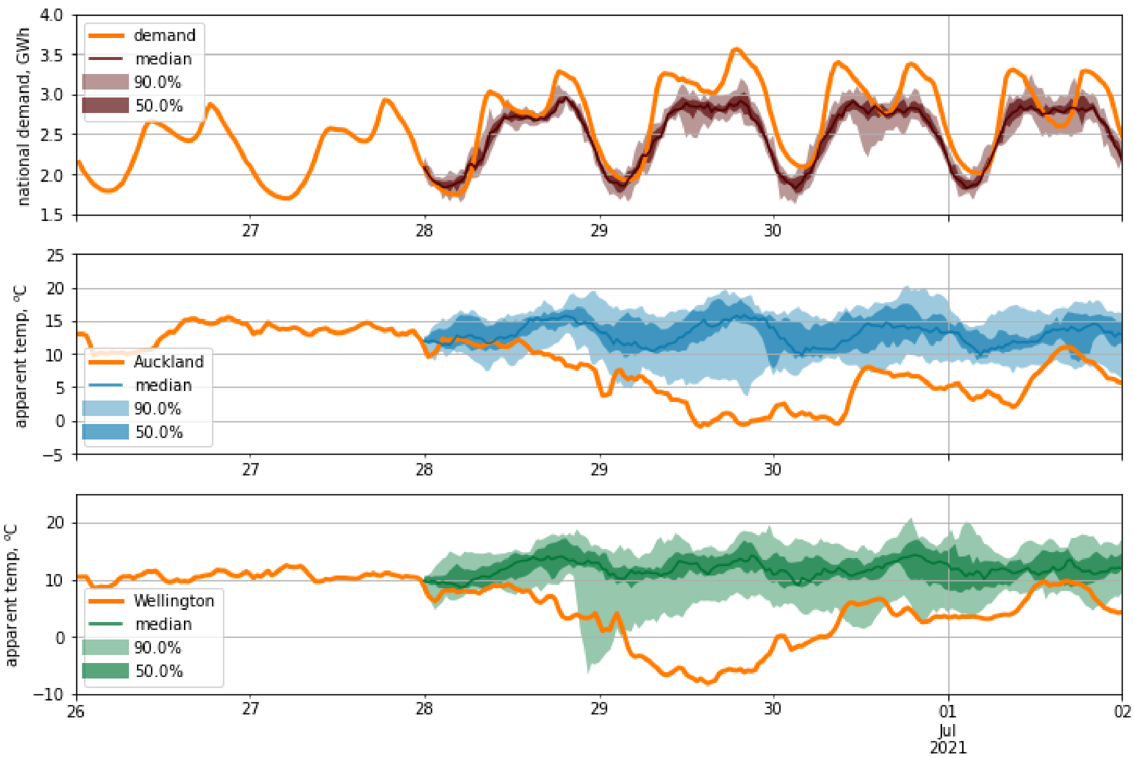

Source: Electricity Authority

One-to-one model

Figure 7: One-to-one model.

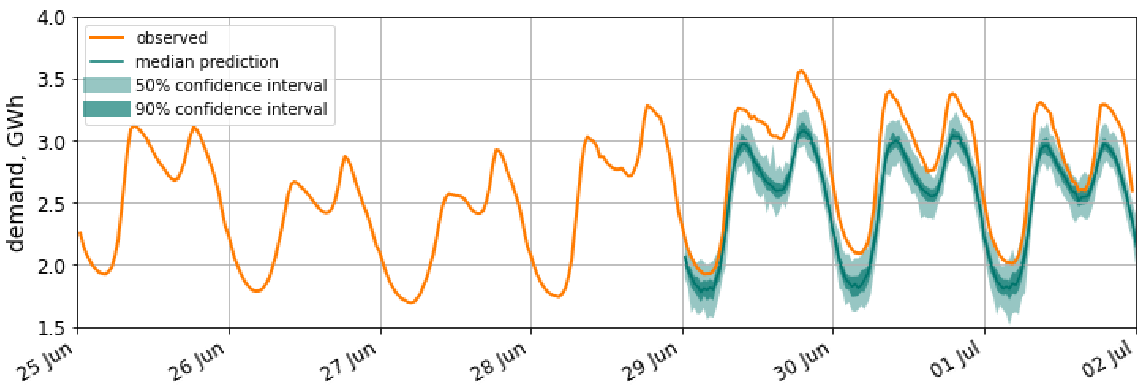

Source: Electricity Authority

2.7 The one-to-one model (Figure 7) only considers demand. The model clearly captures the M-shaped trend over the 48 trading periods, which correspond to the morning and evening demand peaks. The predictions are slightly off since the model does not consider the cold weather.

Many-to-one model, revisited

2.8 In order to work around the issues shown in the many-to-many and one-to-one models, we can consider hourly or half-hourly weather forecasts for some future day. Because it was not possible to obtain the 29 Jun 2021 weather forecasts, we try another set of dates.

2.9 Here we consider the end of Oct 2021. We want to predict future demand for 28-29 Oct 2021. We define 'future demand' as a point in the future when we do not have historical weather data. This means we need to use weather forecasts.

Figure 8: Apparent temperature data for Auckland and Wellington.

Solid lines are observations, dashed lines are forecasts.

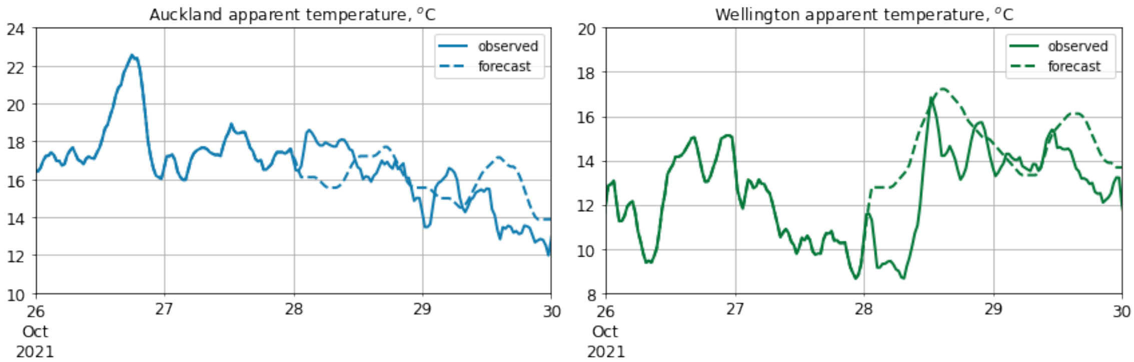

Source: Weather Underground

2.10 Figure 8 shows temperature forecasts (dashed lines) for Auckland and Wellington over 28-29 Oct 2021. The data was acquired on 27 Oct 2021 at roughly 15:00. For comparison, we also plot observed temperatures over these two days (solid lines). This data was acquired after 29 Oct.

2.11 Figure 9 shows the results of our model which is trained using a combination of historical (1 Mar 2014 - 27 Oct 2021) and forecast (28-29 Oct 2021) weather data. The model accurately predicts demand, as there is significant overlap between the confidence intervals and observed demand.

2.12 Even though the weather forecasts do not precisely match the observations, they provide a decent prediction of the overall trend across 28-29 Oct 2021: decrease in Auckland temperature, and steep increase in Wellington temperature followed by gradual decrease.

Figure 9: Many-to-one model.

Uses observed weather data over 1 Mar 2014 to 27 Oct 2021 period to train model. Then uses weather forecasts for 28-29 Oct 2021 to predict demand.

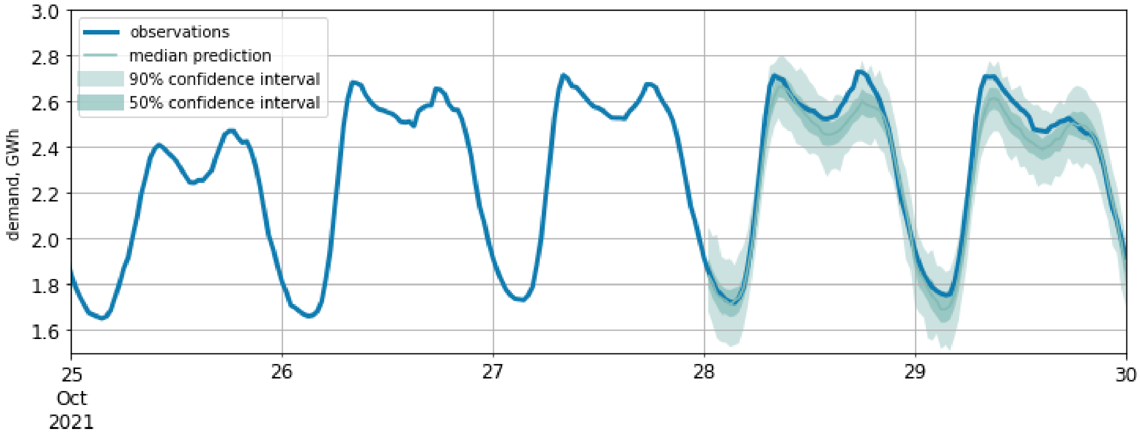

Source: Electricity Authority

2.13 To compare, we now train the model using only historical weather data that covers the 1 Mar 2014 - 29 Oct 2021 period Figure 10. From inspection, both sets of model predictions are highly accurate.

2.14 These results show the importance of using weather data to predict national demand. Compared to the many-to-many model in the above section, training the model to predict only demand is better than asking it to also predict the weather.

2.15 Clearly there can be unpredictable behaviour and large variations in hourly temperature. Hence it is much better to train the model with a combination of historical and forecast weather and use this to predict future demand.

Figure 10: Many-to-one model.

Uses observed weather data for entire 1 Mar 2014 to 29 Oct 2021 period.

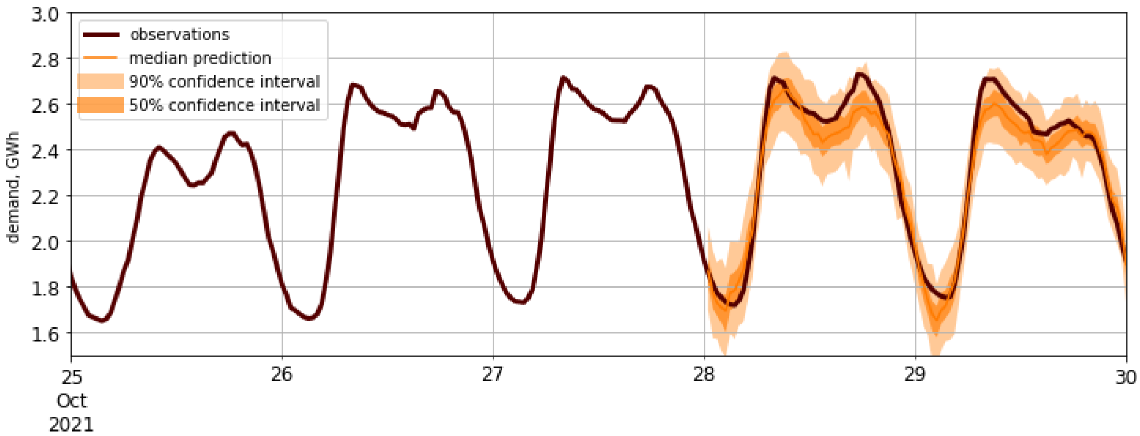

Source: Electricity Authority

2.16 It is possible to numerically evaluate the quality of our forecasts. In GluonTS, we can compute many aggregate performance metrics[9], such as coverage

$$C[\text{quantile}] = \text{mean}\big(Y < \hat{Y}\big), \tag{3}$$

mean absolute percentage error

$$\text{MAPE} = \text{mean}\left(\frac{|Y - \hat{Y}|}{|Y|}\right), \tag{4}$$

mean absolute scaled error

$$\text{MASE} = \frac{\text{mean}|Y - \hat{Y}|}{\text{SE}} \tag{5}$$

with scaled error $\text{SE} = |Y_t - Y_{t-m}|$, and symmetric mean absolute percentage error

$$\text{sMAPE} = 2\ \text{mean}\left(\frac{|Y - \hat{Y}|}{|Y| + |\hat{Y}|}\right) \tag{6}$$

2.17 The metrics are summarised in Table 1.

Table 1: Metrics

Metric	Figure 9	Figure 10
Coverage[0.9]	0.9062	0.8541
MAPE	0.0256	0.0235
MASE	0.5891	0.5553
sMAPE	0.0259	0.0239
Seasonal error	0.1019	0.1019

Source: Electricity Authority

9 R. J. Hyndman and G. Athanasopoulos, *Forecasting: principles and practice*, OTexts (3rd ed.), 2021.