Short-Term Electrical Load Forecasting for Irish Supermarkets with Weather Forecast Data

MSc Research Project
Data Analytics

Martin Orr
Student ID: 19155662

School of Computing
National College of Ireland

Supervisor:
Catherine Mulwa

National College of Ireland
Project Submission Sheet
School of Computing

Student Name:	Martin Orr
Student ID:	19155662
Programme:	Data Analytics
Year:	2021
Module:	MSc Research Project
Supervisor:	Catherine Mulwa
Submission Due Date:	16/08/2021
Project Title:	Short-Term Electrical Load Forecasting for Irish Supermarkets with Weather Forecast Data
Word Count:	5917
Page Count:	20

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.

ALL internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

Signature:	
Date:	22nd September 2021

PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST:

Attach a completed copy of this sheet to each project (including multiple copies).	☐
Attach a Moodle submission receipt of the online project submission, to each project (including multiple copies).	☐
You must ensure that you retain a HARD COPY of the project, both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer.	☐

Assignments that are submitted to the Programme Coordinator office must be placed into the assignment box located outside the office.

Office Use Only	
Signature:	
Date:	
Penalty Applied (if applicable):	

Short-Term Electrical Load Forecasting for Irish Supermarkets with Weather Forecast Data

Martin Orr
19155662

Abstract

With the roll out of smart meters there is great potential to better understand how and when our buildings use energy. This presents an opportunity to rethink how we purchase our energy and also to have a better grasp on our electricity grids. Much of the research in short term electrical load forecasting has been carried out using historic weather data. This works quite well but is not able to account for sudden changes in the weather. This study describes the differences in model results that are found from using both weather forecast and historic data and gives a sense of the value of acquiring weather data. As there can also be sudden level changes in energy load profiles for various reasons it is important that energy prediction tools are able to identify these changes and predict accordingly. This study takes a look at convolutional neural networks as a method of doing this as they have proven effective in identifying shapes in image data. The key findings in this report are that weather forecast data is marginally better than historic data for electricity demand forecasting and that convolutional neural networks are very effective in predicting 30 minute values 24 hours ahead but their predictions for a full day are around the same strength as a linear regression model. There is scope for further research with convolutional neural networks.

1 Introduction

Energy trading produces significant costs for electricity suppliers which then gets passed onto their customers. Approximately 90% of electricity produced is suppliers is purchased by suppliers in the day-ahead market. This results in low liquidity in the intra-day and balancing markets which means that prices are volatile and can be quite high. Accurate day ahead demand forecasting will reduce the amount of electricity that a supplier will have to purchase in the intra-day and balancing markets and thus the supplier will be less exposed to volatile and high prices. More accurate forecasting will improve the efficiency of electrical grid operation, reducing the likelihood of blackouts/brownouts and reduce the need for fossil fuel-burning electricity generation plants running on standby to deal with peak loads and sudden drops in wind energy. This will cut down on greenhouse gas emissions and help Ireland meet their EU-wide emission reduction of 40% over 1990 levels.

1.1 Motivation and Project Background

With the advent of smart meters, the amount of energy consumption data available for analysis has greatly increased. Data Mining and Machine Learning techniques have been proven to be effective in increasing accuracy of electricity demand forecasting. One of the most effective in recent times are models based on the method of deep learning. This method is effective because deep learning can learn about the non-linear relationships that exist between the dependent and independent variables. These relationships are common in electricity consumption due to the sheer number of variables that can impact a building's energy consumption. Another method increasingly leveraged in recent times is Convolutional Neural Networks (CNN). These neural networks can extract features from the shape of the energy load profile from the last 24 hours and make inferences about the building's consumption for the next 24 hours based on this shape. Supermarkets are focused on in this study because over half their electricity consumption is typically due to refrigeration. The amount of energy consumed for refrigeration is highly dependent on weather conditions, particularly the outside air temperature. A failsafe is often triggered in refrigeration plants on hot days which causes the plant's compressors to run at full power without control in order to make sure the food product stays at the correct temperature. This causes energy consumption to remain high until some person sets the refrigeration controls back to automatic mode. This study attempts to identify shape changes in the load profile and the models prediction reflects sudden changes in the energy consumption or forecasted changes in the weather.

1.2 Research Question

Much of the existing research features predictions using historic weather data. This study will attempt to build upon this research by using forecasted weather data from ECMWF and will attempt to answer the following three research questions:

RQ: "To what extent can weather forecast data help improve the accuracy in forecasting the day ahead electricity consumption of a large group of supermarkets in Ireland?

Are convolutional neural networks effective in identifying short term temporal information in the energy load profile?"

Can better accuracy be found when predicting for one building or when predicting for a group of buildings?"

This project provides an in-depth study of the state-of-the-art methods and delivers a solution that can provide an electricity demand forecast for a large group of supermarkets based on tomorrow's weather forecast. The electricity demand forecast is compared to previous studies which used historic weather data and shows the differences in using forecasted weather as a predictor of electricity demand instead of historic weather.

1.3 Research Objectives

Table 1: Objectives and techniques for objective evaluation

obj	Brief Descriptions	Techniques and Evaluation Metric
1	Literature review of recent studies in the areas of forecasting, feature extraction, deep learning and commercial building energy management	N/A
2	Clean and pre-process data. Interpolate for any small gaps in weather data. Discard data where there are large gaps	N/A
3	One-hot encoding to designate time features	N/A
4	Computation of regression analysis and predict energy consumption with historic and forecast weather data	R-Squared, RMLSE
5	Load features into CNN and LSTM and predict energy consumption with historic and forecast weather data	R-Squared, RMLSE
6	Experiment with building level, group supermarkets level	R-Squared, RMLSE

This project is based on trading electricity in the Integrated Single Electricity Market (I-SEM) which Irish electricity suppliers trade in. This report is limited as it is based on weather data from the last 3 years so there may be unseen weather conditions in future that would adversely impact the model's forecast accuracy. The research is limited as well in that it does not account for incorrect weather forecasting. The contribution of this research is the application of CNN models to supermarket data to analyse if it improves prediction. A smaller contribution will be finding the value of using weather forecast data over historic weather data.

This report will critically review the recent literature in the area, followed by a detailed description of the methodology used in implementation, following this the results will be presented, and the report will finish by discussing and concluding the findings.

2 Related Work in Electrical Load Forecasting (2004-2021)

Electricity demand prediction has been a topic of research ever since the first electricity infrastructure was planned. Since then, the time frames for which demand is predicted have expanded to include the next day's electricity demand as electricity traders require this information to buy and sell in the day-ahead markets. Supermarkets have a high portion of their demand, which is made up of refrigeration and Heating, Ventilation and Air Conditioning (HVAC), which are highly dependent on external weather conditions. The main objective of this study is to determine the effectiveness of forecasted weather data at improving the electricity demand predicting model accuracy versus using historical weather data lagged, the latter of which has been much more common in existing research. This literature review will first discuss electricity load forecasting in general. In recent times there have been advances in improving model accuracy through data mining, deep learning and feature selection methods, so these are discussed next. Following this, there is a more in-depth look at the specific methods and nuances that should be accounted for when forecasting short term electricity demand for commercial buildings using weather data. Finally, the literature review findings are summarised, which leads to the next chapter, where how the findings discussed are implemented.

2.1 Electrical Load Forecasting

One review (Nti et al. (2020)) looking at studies in the last decade showed that 90% of studies looked at forecasting energy demand using Artificial Intelligence (AI) while 10% used an engineering and statistics method. Data-driven models have been typically more accurate than deterministic models, but this comes at a cost in that they are difficult to generalise and not transparent (Deb et al. (2017)). Regression models are popular in this field but are better suited to medium and long term forecasting where the significance of periodicity and changes is small (Kuster et al. (2017)). While regression is typically less accurate at short term forecasting, it is worth testing the accuracy with forecasted weather data as this is scarce in the existing literature.

Artificial Neural Networks (ANNs) and time-series methods are more commonly used for short term load forecasting where the consumption patterns are more complex. A 2021 review of data-driven predictive control studies found that many of the studies in the literature had high levels of simplification, having only been applied to single buildings. This study also identified gaps in feature selection, scalability, bench-marking and influence of data quality (Kathirgamanathan et al. (2021)).

In evaluation, 38% and 35% of recent studies used Root Mean Squared Error (RMSE) and Mean Absolute Percentage Error (MAPE), respectively (Nti et al. (2020)). The preferred metric of ASHRAE (American Society of Heating, Refrigerating and Air-Conditioning Engineers) Great Energy Predictor III (2019) was Root Mean Square Logarithmic Error (RMSLE) as it allows for fair comparison of accuracy for models tested on different buildings (Chitalia et al. (2020)). It is worth knowing that this evaluation metric will punish under-prediction more than over-prediction and in that regard it is similar to Root Mean Percentage Square Error (RMPSE)

2.2 Data Mining and Machine Learning Techniques

2.2.1 Deep Learning

Convolutional Neural Networks (CNN), like ANNs mentioned in the previous section, are effective in extracting features with non-linear interactions (Ekundayo (2020)). CNNs have also be leveraged to pick out patterns in the changes in the 1-dimensional electricity consumption profile of a building which are associated with increased energy consumption in the day ahead (Amarasinghe et al. (2017)). A downside of this application is that it can only take information from a time period that is the same size as the convolution layer. Meaning that if the convolutional layer was programmed to take features from the past 24 hours, then something that happened 25 hours ago would be ignored even though it may be affecting the present consumption. Increasing the size of this convolutional layer would reduce computational efficiency and would increase the amount of time required to train the model (Eslami et al. (2020)). An interesting application of CNNs converted multi-variate time series data into images and then used the image processing capability of CNNs to identify features in the data and forecast based on the image features (Sadaei et al. (2019)). Combining CNNs with Long Short Term Memory (LSTM) allows for a model to predict better on the longer-term temporal information in time-series data (Somu et al. (2021)) while the CNN captures the local trend (Tian et al. (2018)). The outputs of the combined CNN and LSTM can be concatenated in a merged module layer, which fuses the features from each. Then the final prediction is made after a fully connected layer. The CNN and LSTM combination has been used to detect theft of electricity and to forecast consumption at residential and grid-level, but the research of this applied to commercial buildings is scarce.

Recurrent Neural Networks (RNN) combined with an attention module improved demand forecasting accuracy by 20-45% against other states of the art models (Chitalia et al. (2020)). The problem of gradient vanishing found in RNNs can be remedied by using LSTM. This maintains the cell state in the RNN and can ignore noisy data commonly found in the consumption profile of supermarkets (Memarzadeh and Keynia (2021), Faisal Mehmood Butt (2021)).

Gated RNNs and CNNs performed better than their non-gated counterparts and the time-series SARIMAX method in a study forecasting instantaneous power load for commercial buildings. This study suggested expanding future work to predict power consumption by load type (Cai et al. (2019)).

2.2.2 Feature Selection

Significant factors should be selected to train the model as this will reduce collinearity and improve computational efficiency (Oprea and Bâra (2019)). There are multiple ways to reduce the number of dimensions fed into the model. Relevance analysis between consumption profile and external factors such as the weather can help recognise the influential factors and associated consumption patterns. One method is to calculate Pearson's product-moment correlation and remove the irrelevant factors. Another method is to conduct Principal Component Analysis (PCA), which can reduce the number of features by combining existing features which has the added benefit of reducing collinearity (Drgoňa et al. (2018)). Clustering of load profiles can be implemented to learn patterns (Jeong et al. (2021)) and analyse the trend in the data (Somu et al. (2021)). Patterns in a previous study were found to last several days, and the history of these patterns determined the consumption in the next day (Jeong et al. (2021)). Curve registration was performed on a grocery store and fed into a logistic mixture Vector Autoregressive Model. By clustering on similarities in consumption patterns, the forecasting error can be reduced (Oprea and Bâra (2019)). This applies to supermarkets in Ireland as there are many different shapes, sizes and compositions, e.g. some will have a café and others will have a clothing section. Synergy, redundancy and the relationship between them can improve feature selection and maximise efficiency (Abedinia et al. (2017)). This study was limited in that only 2-way interactions were considered, so there is potential to extend the algorithm to include higher-order relationships. A feed-forward layer can be added to sparse neural networks to reduce dimensions of input data in an unsupervised way by giving an RNN's inner layers a sparse impression of the information (Gugulothu and Subramanian (2019)).

Wavelet transforms can pre-process the load series by decomposing it into more manageable sub-series, which denote the deterministic and fluctuation parts. This can reduce the irrelevant information fed to the model.(Deb et al. (2017), Rafiei et al. (2018)). Wavelet transformations can also be used to extract discontinuities and breakpoints to remove seasonality (Hošovský et al. (2021)). There is multiple seasonality found in electricity consumption profiles of supermarkets both daily and weekly because of changes in the open/closed status of the supermarket and the annual seasonality, which is caused by changes in the weather.

2.3 Forecasting Energy Demand Of Commercial Buildings With Weather

Regression analysis was effective in forecasting energy consumption for a supermarket in Northern England, and the temperature was found to be more influential on the consumption than humidity (Braun et al. (2014)). The only instance of weather forecast data found in literature multiplied an ASHRAE coefficient by the difference in the high and low forecast temperatures and subtracting this value away from the highest average temperature to predict the average temperature for the next 24 hours (Jia et al. (2016)).

Two studies (Granell et al. (2021) and Rasmussen et al. (2016)) that focused on forecasting energy consumption of supermarkets suggested using different models for operational and non-operational hours to account for how the building will behave differently depending on whether it's open or closed. One-hot-encoding should be used to indicate different timestamps that correspond with holidays, weekends and hours where the supermarket is closed (Nichanian (2020)).

Artificial Neural Networks (ANN) and Support Vector Machines (SVM) don't forecast as effectively as regression models such as ARIMA on linear relationships between variables, and conversely, ARIMA is not effective at the non-linear relationships as ANN combined with ARIMA. An alternative method of dealing with these non-linear relationships is by incorporating a non-linear spline function for both the open and closed hour regimes. While supermarket electricity consumption typically has a linear relationship with temperature, in more modern refrigeration systems where carbon dioxide is the working fluid, the energy consumption can be exponentially higher on hot days.

2.4 Identified Gaps and Conclusions

In conclusion, there have been effective applications of regression analysis, machine learning, time series in predicting short term electricity demand. Although regression is supposedly weaker for this application, it is worth applying it for historical and forecasted weather as it will be clearer to see which has a better correlation. CNNs and LSTMs have shown to be strongest in dealing with non-linear relationships, so these will be trained and tested with forecasted and historic weather also to try to identify the model with the best accuracy in predicting energy. Feature selection will be important as there are many variables in the weather and energy data to ensure computational efficiency and reduce collinearity. Extracting temporal features can be achieved with CNNs, and it is thought that these will be influential on the next day's energy demand in a supermarket. The next chapter will describe the methodology used to carry out these experiments.

3 Short Term Electrical Load Forecast Methodology for Irish Supermarkets

This research has adopted a methodology similar to the Cross-Industry Standard Process for Data Mining (CRISP-DM) methodology. CRISP-DM is a good base to work from because it provides structure to the data mining while still allowing for flexibility as transitions between stages can be reversed. The six stages of the methodology employed for this project can be seen in Figure 1 on the next page.

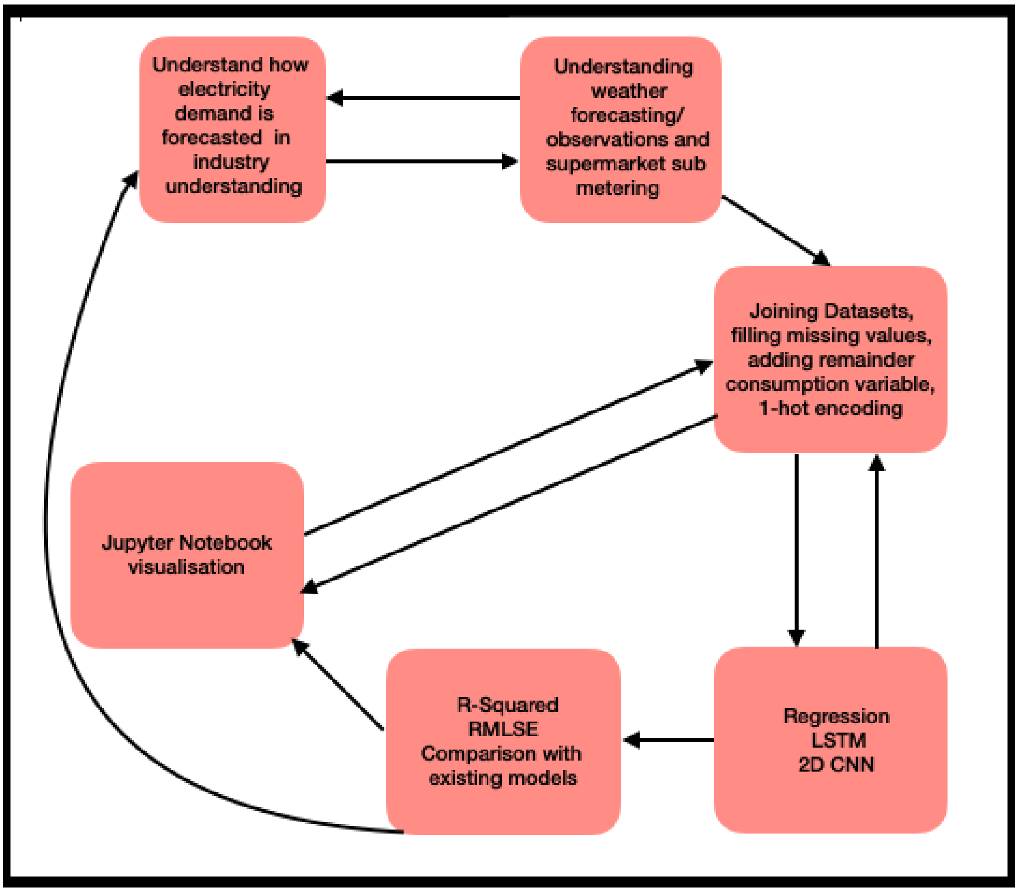

Figure 1: Short Term Electrical Load Forecast Methodology for Irish Supermarkets

3.1 Business Understanding and Problem Understanding

Irish electricity supply companies must predict how much electricity their customers are going to use in the next calendar day each day and purchase this amount of electricity in the day-ahead market. The predictions must be accurate as the electricity supplier will have to buy/sell the difference in the intra-day and balancing market which are much smaller and are subject to high price fluctuation which can leave the electricity supplier financially exposed. It is thought that by dividing the supplier's customers by building use type will be an effective way to group the customers. Summing the predictions of these clusters should lead to a more accurate electricity demand prediction for the supermarket group as a whole and can be an example for how a supplier might group all of their customers.

3.2 Data Understanding

This supermarket is collecting sub-metering data from each of their sites on a half hourly basis. The sub-metering installed at each of the supermarkets has divided the electricity consumption into the different use types (refrigeration, lighting, Heating Ventilation and Air Conditioning (HVAC) etc.) For this study the half-hourly data was queried for refrigeration, HVAC and the main meter. The sum of the refrigeration meter and HVAC meter was subtracted from the main meter to give the remaining consumption. Weather conditions are forecasted 4 times daily for each of the 24 hours ahead by the European Centre for Medium-Range Weather Forecasts (ECMWF) on their Meteorological Archival and Retrieval System (MARS). Historic weather is recorded hourly on ECMWF MARS system as well. From reviewing past research, temperature is the only weather variable that is significant in determining electricity demand in supermarkets so it is the only weather variable selected for this study (Braun et al. (2014)).

3.3 Data Preparation

As the literature review has found, preparation of the data is key in delivering an accurate electricity demand prediction. Pandas, Numpy and Matplotlib are effective Python libraries used to clean the data, perform data manipulations and plot visualisations. Keras & Tensorflow are used as they are user-friendly packages for building deep learning networks. Python will be the main language used for machine learning. The data has been explored on Enacto to investigate the nature of the time series consumption data. Different supermarkets have been grouped based on their size and energy consumption as this makes the computation more efficient and the literature review found that much of the research has focused on predicting energy demand at a building or grid level and that there was a gap in the research for studies carried out on groups of buildings (Kathirgamanathan et al. (2021))

3.4 Modelling

3.4.1 Regression and Time Series

Regression and SARIMA are trained and tested to determine the linear relationship between the weather variables and the next day's electricity demand. This is executed with the daily values of electricity consumption for both historic and forecasted weather data and the accuracies of both are compared.

3.4.2 CNN and LSTM

Refrigeration compressor packs have been known to go into a failsafe "hand" mode where they will run at full power 24/7 rather than the cooling-demand-led "automatic" mode they usually run in. This can be seen in the time series visualisations as a steep step up in consumption that is maintained until a member of the maintenance team switches the compressor pack back to "automatic mode". In order for machine learning modules to learn from changes in the pattern of the time series, a method of identifying these pattern changes must be implemented. The most effective method found in the literature review was CNN. Previous values are interpreted by this as a 2-dimensional image-like dataset.

3.5 Evaluation

The models prediction capabilities were tested on some unseen data and the difference between predicted and actual is the error for each value. The best model is the one with the highest R squared. This is effective to compare models tested on the same data. To compare accuracies of models tested on different data such as finding out whether the models were more accurate on a group of small supermarkets or a group of large supermarkets then a different metric must be used. In this study Root Mean Square Logarithmic Error (RMSLE) is used as this is the metric used in ASHRAE Great Energy Predictor III.

$$RMLSE = \sqrt{\frac{1}{n}\Sigma_{i=1}^{n}(log(p_i + 1) - log(a_i + 1))^2}$$

Where: n = number of values predicted, p = predicted value, and a = actual value

4 Design Specification

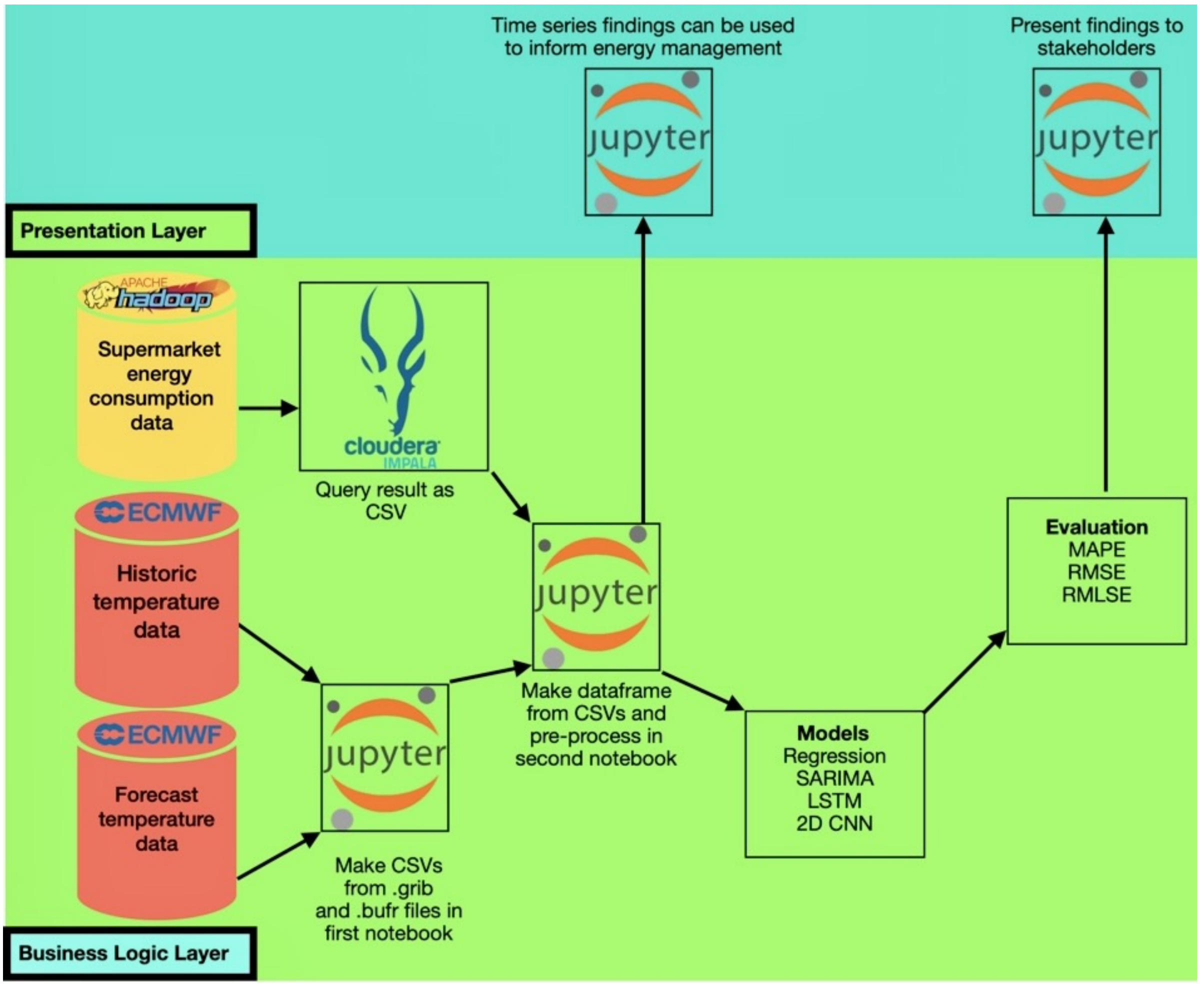

Figure 2: Design Specification

The design for the project is shown in figure 2. Business logic layer is where all the underlying programming, data cleaning and pre-processing, modelling and evaluation sits then the presentation layer sits on top of that.

5 Implementation, Evaluation and Results of Load Forecasting for Irish Supermarkets Models

5.1 Introduction

This chapter discusses the setup and data acquisition, how the prediction models were implemented, how the models were evaluated, and the results from the evaluation of these models.

5.2 Hardware, Experiment Set-up and Data Acquisition

The supermarket's energy consumption data is being uploaded to an energy management platform for visualisations and monitoring. ESB collects this data using an API and loads the data onto their Hadoop database to perform analytics. The data has been downloaded from this using an Impala query as CSV, and this was loaded into Jupyter notebook. The historic and forecasted outside air temperature data were acquired using ECMWF MARS API. The data was pulled for each hour of the day with a separate request for each month. This was found to be the best as larger requests were taking too long and would sometimes never complete. The forecast temperature was originally pulled for the geographic centre of Ireland (53.4°N, 8°W) , but then when the historic temperature was pulled, it was found that the geographic coordinates for both datasets were different because the historic temperature could only be pulled for the coordinates where it was recorded. To account for this, the forecast data pull was rerun to match the average coordinates on the historical temperature data pull (52.57°N, 8.72°W). This rules out geographic reasons for the difference in historic and forecast temperature. The historic temperature data was originally pulled for one location, but upon inspection, it was found that there were many missing data points. Therefore the data pull was rerun to cover a larger area, and this included 2 weather stations. Where there were multiple values for the same hour, the temperature, latitude and longitude were averaged for this data point. The default outputs from the ECMWF MARS system are grib and bufr files (readable in Python by using cfgrib and pdbufr libraries) for forecast and historic, respectively. Each day was compiled into one data frame for all the days, and this was converted to CSV before analysis. These datasets were merged together with the energy dataset as the timestamps were common between all 3.

5.3 Data Extraction and Pre-processing

As noted in chapter 3.1, breaking up the electricity grid into smaller groups should lead to more accurate electricity demand predictions. This research looks at 11 supermarkets as a group, and 1 supermarket on its own and the results are compared. As supermarkets have different consumption patterns on different days of the week because of the difference in opening hours, one-hot encoding has been applied to the data to denote the different days. The historical weather data has been lagged forward by 24 hours, so it matches with the previous day's values. This was the method found most commonly in the literature and allows for the best comparison of predictions using historical weather data and forecasted weather data. Outliers were identified as being far from the main group on the frequency distribution and changed to be the median (calculated, not including outliers). Data is cleaned by removing duplicate values and interpolating weather data where there are small gaps (¡ 4 hours). Where there are gaps in the mains consumption energy data, these values have been removed from the model training as a target as it is impossible to estimate what the mains consumption would have been within these gaps, and it would lead to a poorly trained model. To account for these gaps in the input data to the model, the number of main meter consumption values has been queried for each day. Where the day doesn't have the full amount expected in mains meter values, then this day was deleted. This wouldn't have affected the model too much, as the largest gap was 2 days. The data was re-sampled to sum the values each day for regression and SARIMA models.

5.4 Data Exploration

Time series plots were made for each of the variables to identify seasonality. Data was explored using Python's quickda library. Distribution plots were effective to show the density of consumption values and skewness/Kurtosis. Correlation matrix plots showed a correlation between the variables.

	mains_consumption	hvac_consumption	refrigeration_consumption	remainder_consumption	forecast_temp	historic_temp
mains_consumption	1.00	0.66	0.59	0.92	0.22	0.36
hvac_consumption	0.66	1.00	0.06	0.67	-0.13	-0.04
refrigeration_consumption	0.59	0.06	1.00	0.24	0.50	0.57
remainder_consumption	0.92	0.67	0.24	1.00	0.05	0.19
forecast_temp	0.22	-0.13	0.50	0.05	1.00	0.87
historic_temp	0.36	-0.04	0.57	0.19	0.87	1.00

Figure 3: Correlation Matrix for the initial combined dataset for a group of supermarkets

	mains_consumption	hvac_consumption	refrigeration_consumption	remainder_consumption	forecast_temp	historic_temp
mains_consumption	1.00	0.27	0.67	0.90	0.19	0.32
hvac_consumption	0.27	1.00	-0.23	0.33	-0.20	-0.22
refrigeration_consumption	0.67	-0.23	1.00	0.31	0.44	0.53
remainder_consumption	0.90	0.33	0.31	1.00	0.01	0.14
forecast_temp	0.19	-0.20	0.44	0.01	1.00	0.87
historic_temp	0.32	-0.22	0.53	0.14	0.87	1.00

Figure 4: Correlation Matrix for the initial combined dataset for a single supermarket with a high frequency of refrigeration compressor packs going into manual mode

The figures, 3 and 4 show that there is a lower correlation between refrigeration consumption and temperature in the single store than in the group of stores. This may be because the compressor packs in this store were running on full power in manual mode, whereas the compressor packs in the group of supermarkets were automatically controlled and would have varied the compressor power up and down depending on the outside air temperature.

5.5 Implementation, Evaluation and Results of Regression

The resampled daily data was used to carry out some regression analysis.

5.5.1 Implementation

The target and weather forecast variables were shifted back a day to simulate the model predicting the day ahead with the variables from the day before. All the variables apart from the target and the historic temperature were set as input to the forecast regression, and all the variables apart from the target and forecast temperature were set as the input to the historic regression. A train:test split of 8:2 was placed on both sets. Both models were fit on both sets of training data then used to predict on the test data

5.5.2 Evaluation

These models were evaluated using an R squared formula which compared the real target variable and the target variable predicted by the linear model. This was completed for both only the unseen data and for all the data. The RMLSE was also calculated using the SciKit-Learn function

Table 2: R squared and RMLSE for linear regression

Model	**Dataset**	**R-Squared**	**RMLSE**
Forecast	Group of 11 large Supermarkets	0.371	0.03578
Historic	Group of 11 large Supermarkets	0.415	0.03436
Forecast	Single Supermarket	0.462	0.0457
Historic	Single Supermarket	0.391	0.0486

Table 2 shows the evaluation metrics for linear regression. It is surprising to find that forecast data provided a better R squared for the single supermarket but a worse R squared value for the group of supermarkets.

Figure 5: Regression predictions: forecast vs historic

Figure 5 shows the predictions of both linear regression models against what the actual values were.

5.6 Pre-Implementation 2D CNN and LSTM

Before beginning LSTM and 2D CNN, the forecast temperature data was shifted up by 48 rows (1 day) to be used as the predictor for the day ahead. The target variable that was being predicted was 48 half hours ahead of the data being fed to the model. This simulates the model having to predict the half-hour consumption value 24 hours in advance using past weather data or future forecast data. A sliding window view of the 11 variables (3 energy consumption, 1 weather variable and 7 1-hot encoded days of the week) to achieve the desired input for the CNN. For each target, there was an 11x48 array of input values. A model each for the forecast temperature and the historic temperature was created using LSTM and 2D CNNs [1]. The model from this footnote was modified to remove the target variable from the training of the dataset because the target variable here had been shifted back by a day. There was no benefit from including the target variable in the training as the target variable was the sum of 3 of the features (refrigeration consumption, HVAC consumption and remainder consumption). In these models the previous values being fed into the model are windowed, and this window is what the model uses to predict the target variable. Each variable was split out into the train and test parts with a ratio of 0.8:0.2 respectively. Then the train and test parts were reshaped to an array of length of the split and width of the window. The array for each feature was concatenated into a large array.

A random number was seeded so that the CNNs would be randomised with the same weights, and the same results could be achieved each time. Values were scaled so zero was minimum and 1 was maximum. The half-hourly consumption predictions for both historic and forecast temperature were resampled to get the daily mains consumption value. These values were compared with the real mains consumptions for that day to calculate RMSE and RMLSE.

5.7 LSTM

5.7.1 Implementation

The LSTM was trained using historic and forecast datasets for 5 epochs with shuffle set to false.

5.7.2 Evaluation

After the model was trained it was used to predict the mains consumption in the unseen test data. The R-squared for these predictions was calculated using the r2_score from scikit-learn python library. The values for this and the RMLSE for LSTM can be seen in Table 3.

[1] https://towardsdatascience.com/time-series-forecasting-with-2d-convolutions-4f1a0f33dff6

Table 3: R squared and RMLSE for LSTM

Model	Dataset	R-Squared	RMLSE
Forecast LSTM	11 large Supermarkets	-8.312	0.122
Historic LSTM	11 large Supermarkets	-1.657	0.0710
Forecast LSTM	Single Supermarket	-0.130	0.0497
Historic LSTM	Single Supermarket	-2.300	0.0804

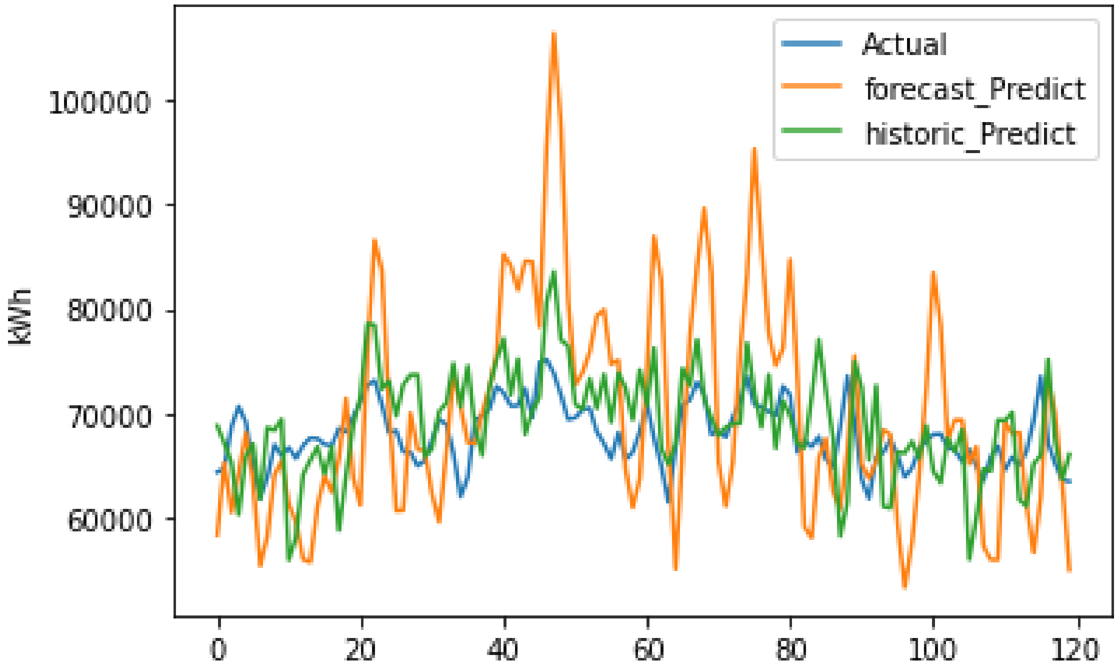

Figure 6: LSTM Daily: forecast vs historic

The values for forecast LSTM predictions seen in Figure 6 can be described as erratic. There are large negative values for R squared in this dataset which is put down to a calculation error as from the visualisation it looks like there could be positive correlation.

5.8 2D CNN

Firstly a 1 dimensional CNN was made that went through each feature in a row. After this, 2 CNNs were built that would look over the data like a 2D image. Different window sizes were experimented with, and it was found that 48 produced the least error. The number of filters (the number of neurons in the first row) was selected as 24.

5.8.1 Implementation

A random number was seeded at the start so that the CNNs would be randomised with the same weights, and the same results could be achieved each time. Different sizes (5, 10, 20, 40) were experimented with for the filter size in the 2D CNN with little differences in the error, so a size of 20 was selected as this allowed for relatively fast computing time. Different window sizes of data were selected. Any lower than 48 would not work with the shape of the training data. Bigger windows meant longer time to train the model. The best performing CNN layout can be seen in Figure 7

```
Layer (type)                 Output Shape              Param #
=================================================================
module_wrapper_295 (ModuleWr (None, 48, 11, 24)        72

_________________________________________________________________
module_wrapper_296 (ModuleWr (None, 12672)             0

_________________________________________________________________
module_wrapper_297 (ModuleWr (None, 1000)              12673000

_________________________________________________________________
module_wrapper_298 (ModuleWr (None, 100)               100100

_________________________________________________________________
module_wrapper_299 (ModuleWr (None, 1)                 101
=================================================================
Total params: 12,773,273
Trainable params: 12,773,273
Non-trainable params: 0
_________________________________________________________________
```

Figure 7: Layout of CNN: 1xConv layer, 1xFlatten layer followed by 3xDense layers

5.8.2 Evaluation

The half-hourly consumption predictions for both historic and forecast temperature were resampled to get the daily mains consumption value. These values were compared with the real mains consumptions for that day to calculate R-squared and RMLSE. The best CNN model had a data size of 48 half hours (1 day), it had a sequential architecture of 2x2D convolution layers, followed by a flatten layer, followed by 3 dense layers. As the dataset was divided 8:2 it meant that the test set began in the middle of the day. To remedy this for evaluation the dataframe was reversed as it was known that the last value was at the end of the day. From here the values were summed in groups of 48 and the remainder was deleted. The dataframe was then re-reversed so that the visualisation would match that of other models.

Table 4: R squared and RMLSE for 2D CNN

Model	**Dataset**	**R-Squared**	**RMLSE**
Forecast 1 Conv Layer	11 large Supermarkets	0.374	0.0346
Historic 1 Conv Layer	11 large Supermarkets	0.377	0.0345
Forecast 1 Conv Layer	Single Supermarket	-2.155	0.0824
Historic 1 Conv Layer	Single Supermarket	-5.935	0.1188
Forecast 2 Conv Layer	11 large Supermarkets	-1.534	0.0679
Historic 2 Conv Layer	11 large Supermarkets	0.067	0.0408
Forecast 2 Conv Layer	Single Supermarket	-0.256	0.0526
Historic 2 Conv Layer	Single Supermarket	-0.356	0.5474

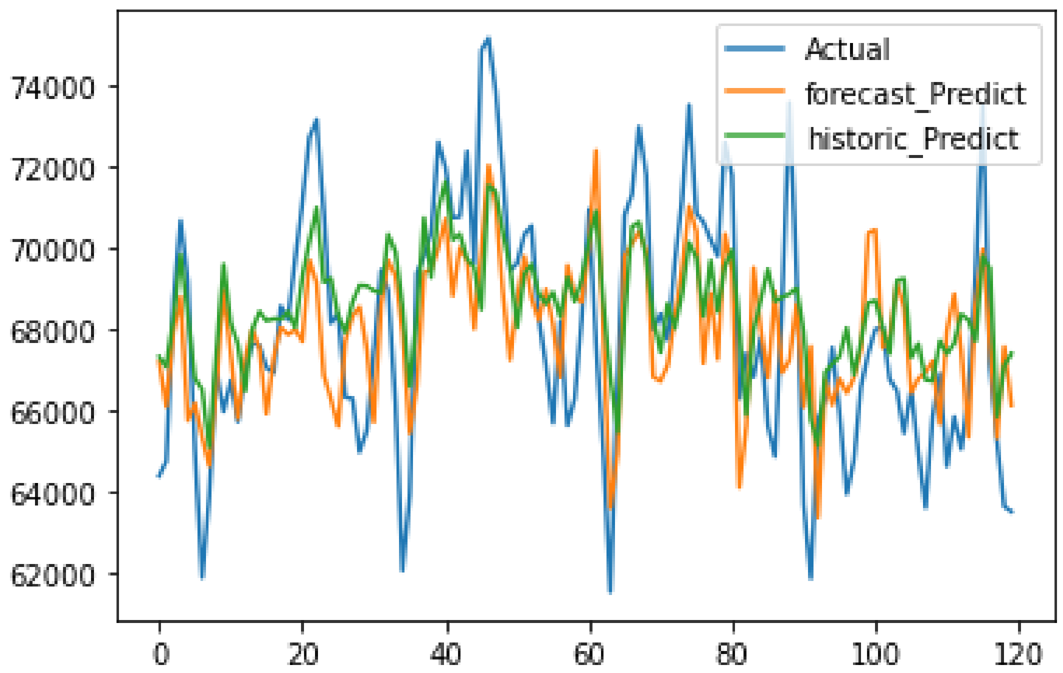

Figure 8: CNN 1xConv Layer daily predictions: forecast vs historic

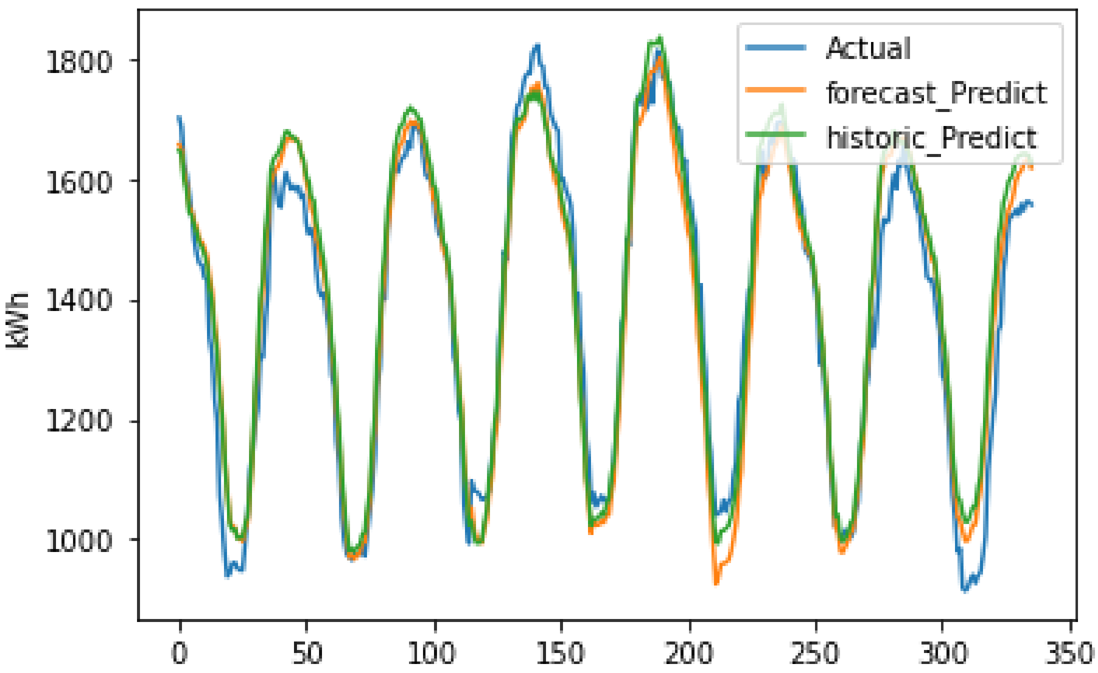

Figure 9: CNN 1xConv Layer half-hour predictions: forecast vs historic

The model with the best half-hourly predictions can be seen in Figures 8 and 9. This model had a half-hourly R squared value of 0.94 (RMLSE = 0.0467) and 0.937 (RMLSE =0.04777) for forecast and historic, respectively, but the daily R squared was just 0.37.

6 Discussion and Comparison

Interestingly, the forecasted temperature was 0.6°C higher than the historic temperature for both data sets.

6.1 Comparison of Developed Models

When the window size was increased to 3 days (1 day provided the most accurate day ahead results) for CNN and LSTM, results showed that the accuracy was better for historic weather data than forecasted. This discrepancy may be down to errors in the weather forecast.

Both linear regression and CNN 2D had a lower RMLSE and R-Squared with forecast data instead of historic data. It was surprising to find that the CNN has a higher R squared value and also had a higher RMLSE than the linear regression. Linear regression is the model commonly used in the energy industry to forecast energy. It is promising to see that one of the CNN models was of a similar R Squared value to the linear regression models.

In the next and final chapter the findings are concluded with respect to the research questions and some ideas for future work are recommended.

7 Conclusion and Future Work

In conclusion, it can be said that using forecast data will lead to more accurate predictions of day-ahead electricity demand for supermarkets, although the difference is very small. Only one of the CNNs and LSTMs had an accuracy comparable with Linear Regression. This was the 1 layer CNN and it had an R Squared of 0.37 which is the same as the linear regression on the same dataset. 2D CNN predicted quite well for half an hour, 1 day ahead with an R squared of 0.94, but when these half hours were summed up it just had a value of 0.39 correlation. There may have been an error in the method calculating the forecast for the full day. R Squared and error values were subject to change drastically depending on what time range of data was fed to the models. This suggests that there is a lot of variability in the target variable making it hard to predict. The reason for there not being that much research featuring groups of buildings could be because people are unwilling to take the plunge. This study shows that models trained on groups of buildings yield less accuracy than single buildings so it may be harder to get credit for a good model. It was expected that better R squared values would be found in the single supermarket dataset but due to what is believed to be calculation errors this was not the case. With a half hour R squared value of 0.94 for the day ahead it can be said that CNNs are effective in dealing with the level changes that make up part of the electricity load profile for a supermarket and should be effective at detecting refrigeration packs that have gone into manual mode but further detailed research is needed in this area.

7.1 Future Recommendations

Future studies in this area should incorporate the gas consumption in the model as this would affect the HVAC and refrigeration consumption. Gas consumption data was not available for this study. Using an optimizer would help improve the CNN outputs and the RMLSE optimiser would be useful for models competing in the ASHRAE great energy predictor. Building multiple CNNs for different parts of the day is a good direction for future research as the method in this study was forecasting for a half hour period 24 hours ahead. One conclusion and recommendation for future research using ECMWF data is that there is no need to use Linux to access the MARS request API. It can be done natively on a UNIX system on the M1 chip system using MiniForge and Jupyter notebook. Future research could investigate building a classifier with alarms for when the refrigeration pack goes into manual mode by analysing the consumption data.

Acknowledgements I'd like to thank NCI for facilitating this research and especially my supervisor, Catherine Mulwa. Her regular meetings and helpful advice are an exemplar of what a great thesis supervisor should be. Thanks to my employer, ESB, and my manager, Tom Mooney, for guidance and providing study leave when I needed it. Thanks to Tesco for providing me with their data so that this research could be done. Thanks to ECMWF and the people at Met Eireann who kindly set me up with access to the API so that I could access forecast and historic data. Last but not least I'd like to thank my girlfriend, parents and friends for their moral support.

References

Abedinia, O., Amjady, N. and Zareipour, H. (2017). A new feature selection technique for load and price forecast of electrical power systems, *IEEE Transactions on Power Systems* **32**(1): 62–74.

Amarasinghe, K., Marino, D. L. and Manic, M. (2017). Deep neural networks for energy load forecasting, *2017 IEEE 26th International Symposium on Industrial Electronics (ISIE)*, pp. 1483–1488.

Braun, M., Altan, H. and Beck, S. (2014). Using regression analysis to predict the future energy consumption of a supermarket in the uk, *Applied Energy* **130**: 305–313.
URL: *https://www.sciencedirect.com/science/article/pii/S0306261914005674*

Cai, M., Pipattanasomporn, M. and Rahman, S. (2019). Day-ahead building-level load forecasts using deep learning vs. traditional time-series techniques, *Applied Energy* **236**: 1078–1088.
URL: *https://www.sciencedirect.com/science/article/pii/S0306261918318609*

Chitalia, G., Pipattanasomporn, M., Garg, V. and Rahman, S. (2020). Robust short-term electrical load forecasting framework for commercial buildings using deep recurrent neural networks, *Applied Energy* **278**: 115410.
URL: *https://www.sciencedirect.com/science/article/pii/S0306261920309223*

Deb, C., Zhang, F., Yang, J., Lee, S. E. and Shah, K. W. (2017). A review on time series forecasting techniques for building energy consumption, *Renewable and Sustainable Energy Reviews* **74**: 902–924.
URL: *https://www.sciencedirect.com/science/article/pii/S1364032117303155*

Drgoňa, J., Picard, D., Kvasnica, M. and Helsen, L. (2018). Approximate model predictive building control via machine learning, *Applied Energy* **218**: 199–216.
URL: *https://www.sciencedirect.com/science/article/pii/S0306261918302903*

Ekundayo, I. (2020). *Optuna optimization based cnn-lstm model for predicting electric power consumption*, Master's thesis, Dublin, National College of Ireland.
URL: *http://norma.ncirl.ie/4440/*

Eslami, E., Choi, Y., Lops, Y., Sayeed, A. and Salman, A. K. (2020). Using wavelet transform and dynamic time warping to identify the limitations of the cnn model as an air quality forecasting system, *Geoscientific Model Development* **13**(12): 6237–6251.
URL: *https://gmd.copernicus.org/articles/13/6237/2020/*

Faisal Mehmood Butt, Lal Hussain, A. M. (2021). Artificial intelligence based accurately load forecasting system to forecast short and medium-term load demands, *Mathematical Biosciences and Engineering* **18**: 400–425.
URL: *https://www.aimspress.com/article/10.3934/mbe.2021022*

Granell, R., Axon, C. J., Kolokotroni, M. and Wallom, D. C. (2021). Predicting electricity demand profiles of new supermarkets using machine learning, *Energy and Buildings* **234**: 110635.
URL: *https://www.sciencedirect.com/science/article/pii/S0378778820334216*

Gugulothu, N. and Subramanian, E. (2019). Load forecasting in energy markets: An approach using sparse neural networks, *Proceedings of the Tenth ACM International Conference on Future Energy Systems*, e-Energy '19, Association for Computing Machinery, New York, NY, USA, p. 403–405.
URL: *https://doi.org/10.1145/3307772.3330167*

Hošovský, A., Piteľ, J., Adámek, M., Mižáková, J. and Židek, K. (2021). Comparative study of week-ahead forecasting of daily gas consumption in buildings using regression arma/sarma and genetic-algorithm-optimized regression wavelet neural network models, *Journal of Building Engineering* **34**.
URL: *https://www.scopus.com/inward/record.uri?eid=2-s2.0-85096400935&doi=10.1016%2fj.jobe.2020.101955&partnerID=40&md5=58e61beb2bccdea55338f1c2216f8c63*

Jeong, D., Park, C. and Ko, Y. M. (2021). Short-term electric load forecasting for buildings using logistic mixture vector autoregressive model with curve registration, *Applied Energy* **282**: 116249.
URL: *https://www.sciencedirect.com/science/article/pii/S0306261920316408*

Jia, J., Xing, J., Ling, J. and Peng, R. (2016). A method to predict cooling load of large commercial buildings based on weather forecast and internal occupancy, *Frontiers in Energy* **10**(4): 459.

Kathirgamanathan, A., De Rosa, M., Mangina, E. and Finn, D. P. (2021). Data-driven predictive control for unlocking building energy flexibility: A review, *Renewable and Sustainable Energy Reviews* **135**: 110120.
URL: *https://www.sciencedirect.com/science/article/pii/S1364032120304111*

Kuster, C., Rezgui, Y. and Mourshed, M. (2017). Electrical load forecasting models: A critical systematic review, *Sustainable Cities and Society* **35**: 257–270.
URL: *https://www.sciencedirect.com/science/article/pii/S2210670717305899*

Memarzadeh, G. and Keynia, F. (2021). Short-term electricity load and price forecasting by a new optimal lstm-nn based prediction algorithm, *Electric Power Systems Research* **192**: 106995.
URL: *https://www.sciencedirect.com/science/article/pii/S0378779620307938*

Nichanian, S. (2020). *Understanding the impact of covid-19 on electrical demand*, Master's thesis, Dublin, National College of Ireland.
URL: *http://norma.ncirl.ie/4458/*

Nti, I. k., Adekoya, A., Nyarko-Boateng, O. and Teimah, M. (2020). Electricity load forecasting: a systematic review, *Journal of Electrical Systems and Information Technology* **7**.

Oprea, S. and Bâra, A. (2019). Machine learning algorithms for short-term load forecast in residential buildings using smart meters, sensors and big data solutions, *IEEE Access* **7**: 177874–177889.

Rafiei, M., Niknam, T., Aghaei, J., Shafie-Khah, M. and Catalão, J. P. S. (2018). Probabilistic load forecasting using an improved wavelet neural network trained by generalized extreme learning machine, *IEEE Transactions on Smart Grid* **9**(6): 6961–6971.

Rasmussen, L. B., Bacher, P., Madsen, H., Nielsen, H. A., Heerup, C. and Green, T. (2016). Load forecasting of supermarket refrigeration, *Applied Energy* **163**: 32–40.
URL: *https://www.sciencedirect.com/science/article/pii/S0306261915012738*

Sadaei, H. J., de Lima e Silva, P. C., Guimarães, F. G. and Lee, M. H. (2019). Short-term load forecasting by using a combined method of convolutional neural networks and fuzzy time series, *Energy* **175**: 365–377.
URL: *https://www.sciencedirect.com/science/article/pii/S0360544219304852*

Somu, N., Raman M R, G. and Ramamritham, K. (2021). A deep learning framework for building energy consumption forecast, *Renewable and Sustainable Energy Reviews* **137**: 110591.
URL: *https://www.sciencedirect.com/science/article/pii/S1364032120308753*

Tian, C., Ma, J., Zhang, C. and Zhan, P. (2018). A deep neural network model for short-term load forecast based on long short-term memory network and convolutional neural network, *Energies* **11**: 3493.