Investigation of Active Learning Techniques For dynamic Time-of-Use (dToU) Tariff Policy Design for Residential Users

by

S. Nikte

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended on Friday September 20, 2019 at 09:00 AM.

Student number:	4716132	
Project duration:	January 10, 2019 – September 20, 2019	
Thesis committee:	Prof. Dr. P. Palensky,	TU Delft,committee chair
	Dr. S. Tindemans,	TU Delft, supervisor
	Dr. M. Spaan,	TU Delft, committee member

TU Delft

Acknowledgement

I would first like to thank my thesis advisor Dr. Simon Tindemans of the Intelligent Electric Power Grid (IEPG) research group at Delft University of Technology. I got all the help from Prof. Tindemans' whenever I ran into a trouble or had a question about my research or writing. I am thankful for his honest comments during the period of the thesis. I truly enjoyed the process brainstorming and development of new methods and experiments.

I would also like to acknowledge Yashaswi Giridhar and Devendra Kulkarni as the second readers of this thesis, and I am gratefully indebted for their very valuable comments on this thesis.

I would also like to thank my dear friends in Delft who made this city my second home. From serious intellectual debates to some fun teasing, they helped me refresh and relax during my studies. They fed me through my sleepless nights of work, helped me through toughest courses, enjoyed music with me, or simply offered me tea whenever I met them. The list is long and I would probably miss someone, but here we go- Thank you Chinmay, Yashaswi, Devendra, Sneha, Rajesh, Subranshu, Girish, Hsoc, Shardul, Biswadeep, Sukanya, Husain, Pranav, Karishma, Gaurav and Rohit for sharing last two years together.

Finally, I must express my very profound gratitude to my parents for providing me with continuous support and supporting me in every endeavour that I chose to follow. I really appreciate them for believing in me and giving me freedom to pursue my passion. This accomplishment would not have been possible without them. This thesis is dedicated to them.

S. Nikte
Delft, September 2019

Abstract

The active learning approach is a special case of semi-supervised machine learning which is able to interactively query the user to reduce the uncertainty of the machine learning model. The approach is useful to minimize the data labelling cost. The project aims to study and use this method to characterize residential electricity users' demand response to improve the prediction accuracy of energy demand. During the trials, the policies which penalise or incentivize the users to change their behaviour involve a cost associated with the grid management. Therefore, the experiments which include the above-mentioned policies are considered as cost bearing experiments. The goal of this project is to study the effect of selective sampling and random sampling of such cost bearing experiments on energy consumption prediction accuracy in simulated residential energy consumption environment. Firstly, we show the simulator design for simulating demand response of users under dynamic tariff policy. Then, we investigate two selective sampling methods- variance reduction and novelty detection. The performance of these methods under various criteria is analysed.

List of Figures

List of Tables

Contents

1

Introduction

1.1. Background

Forecasting of residential users electricity consumption is a critical aspect of electricity retailer's work. The function of electricity retailing is to sell electricity to end-customer. Electricity retailing is the fourth step after generation, transmission and distribution of electricity.

In traditionally regulated electricity markets, electricity retailers usually have vertical structure i.e. they are also involved in generation, transmission and distribution of electricity. In this case, the forecasting of user demand helps the retailer to set the generation of electricity accordingly. In deregulated electricity markets, an electricity retailer is usually an independent entity which buys the electricity from the wholesale market and sells it in the retail market. An electricity retailer has to bid for electricity in advance. The forecast of electricity consumption of their users may help them evaluate their needs and therefore, buy adequate electricity from the electricity generators.

Demand response is a tool to influence and alter the consumption pattern of end-users to match the generation of electricity. Demand response schemes usually provide users with financial incentives to alter their electricity consumption pattern. The users are usually guided with signals from an electricity retailer. Dynamic tariff policies are one of such demand response schemes. In this scheme, the users are charged with different tariff prices for their electricity consumption for different periods of a day.

The electricity consumption of users under dynamic tariff signals may change from their usual consumption levels. Therefore, electricity retailers need to design 'experimental' demand response schemes to understand user electricity behaviour pattern under dynamic tariff schemes. The experiment in such case can be a 'dummy' dynamic tariff signal to observe the user response. When these experiments are performed under a variety of conditions, they provide information about user responses under those conditions. As in later chapters, we will know, these experiments are costly. Artificially changing user behaviour by 'dummy' dynamic tariff involves the cost of electricity grid management, balancing fees [41] etc. Implementation of such experiment trials requires a lot of planning and careful execution. Therefore, we cannot perform these experiments continuously.

This thesis deals with the following problem - how to improve the performance of the forecasting model for electricity demand under dynamic tariff when the number of dynamic tariff signals is constrained. The following section defines the problem that we will handle in this research project.

1.2. Problem definition

Weather conditions and seasonal habits of the users are important predictors of the usual electricity consumption of residential user[43]. But unlike large industrial users, the residential users might change their usual demand response behaviour for the sake of convenience[12]. The very common example of such case is switching on the TV for an important broadcasting event. Although, individual television is an insignificant load when considering the consumption of large cluster of users, the combined effect of such actions by a large number of users has a significant effect on the overall consumption of such cluster. This consumption behaviour thus evades any determinate prediction model and as a result, exhibits a stochastic quality.

To understand the consumption behaviour of residential users (which also changes based on geographical location, local culture and habits) the utility companies need to run pilot projects to carry out experiments to understand the user demand response for various situations. We restrict ourselves to system behaviours such as that of dynamic tariff signals as a means to analyze user response as a function of changing financial incentives during a day. As these experiments can potentially affect the operations and stability of entire electricity grid, planning and implementation of such pilot projects place considerable stress on the resource availability - both financial as well as human - on an organization. Therefore, such projects are bound by the availability of resources and time. One also has to consider the high cost of hardware per kW-Hr (kilo Watt Hour) consumption of electricity. Unlike industrial users, the residential loads have a larger user base (number of users) but smaller individual consumption. As a means to navigate these challenges, it becomes necessary for the utility companies to improve the quality of each experiment conducted under the pilot project.

1.3. Problem statement

In this work, the author aims at analyzing and providing a mechanism for the easing of constraints faced by the utility companies to process residential user demand response for dynamic tariff signals. The problem of forecasting of demand response is constrained by the availability of the various resources including human (experts from industry), material (required hardware and software for gathering the data) and money (budget). It is assumed that the cost of each dynamic tariff experiment is equal and does not need to be quantified in financial units.

Additionally, this research project presents a dynamic tariff response simulator for dealing with the limited data of original dataset. Overall, the project captures the following main research questions linked to the problem and the objective of the work:

- What are the methods used in demand response?
- What is the contribution of various weather conditions and calendar effects to the forecasting of residential user electricity consumption?
- How does a simple regression forecasting model (baseline model) perform under a random selection of experiment days?
- How to improve the selection of experiment days such that information gained from the experiments would be maximized?
- How does the model with the selective sampling of experiment days perform compared to the baseline model? How to quantify the performance in terms of budget?

Along with the reported research questions, the research project also focuses on the design and development of a dynamic tariff demand response simulator due to the limited availability of the data.

- How to develop a dynamic tariff demand response simulator when limited data is available?

1.4. Research aim and contributions

While considering the overview of the project, the research aim and contribution in this project is two-fold.

- The primary aim of the project is to develop a novel method for the selection of potential information-rich experiment days to improve the accuracy of the overall demand response forecast. We try to find algorithmic solutions to improve the existing forecasting models of demand response of user consumption for a limited number of dynamic tariff experiments. In other words, to improve the accuracy of the forecasting, we aim to intelligently select the experimental conditions which would provide greater insights into user consumption habits.

  The method of selective sampling of data label is known as active learning. Compared to regression settings, more literature is available for the classification problems of active learning framework[37]. however, the forecasting of a demand response consumption is a regression problem, which limits the available literature for this research work. We investigate the two selective sampling methods for selecting the experiment days.

- The secondary aim of this project is to develop a dynamic tariff demand response simulator to tackle the problem of limited availability of the data. During the tenure of this research project, unavailability of adequate data of dynamic tariff demand response has led to the development of the simulator of the experiments. This simulator simulates the fixed tariff response and dynamic tariff response of users for generated tariff policy. The project has used the dataset from 'Low Carbon London' (LCL) project [24] which collected the demand response consumption for a dynamic Time-of-Use (dToU) tariff policies. The data-driven simulator model is based on the LCL dataset and it is designed based on the analysis and findings of the LCL project.

1.5. Thesis outline

The contents of the following chapters are discussed below. Each chapter follows a specific layout: introduction and background, the main body of the topic, a brief conclusion summarizing the discussion and the discussion addressing the research questions listed above.

Chapter 2 discusses the various available demand response methodologies and their effect on the electricity grid. It also addresses the problem of residential demand response from the perspective of a utility company or an electricity distributor. The chapter shows the effect of demand response for the stability and uninterrupted operation of an electricity grid.

Chapter 3 dives into the domain of machine learning. It provides a brief introduction of machine learning, followed by the introduction of decision trees and random forest models. It also describes the one-class SVM models considered for the research project. Then, the chapter provides an introduction to the three important active learning frameworks, namely, membership query synthesis, stream-based selective sampling and pool-based selective sampling. The various methods of query formation are then discussed followed by the discussion about the useful methods.

Chapter 4 gives a gentle introduction to the problem setup and design by introducing various elements of the energy retail market and the Time-of-Use tariff policy. This chapter is a recommended read for the readers who are not familiar with the electricity distribution domain. This chapter also provides the assumptions considered in the design of the setup of this research project.

Chapter 5 provides details about the data used for the research project. The information about the fixed tariff consumption dataset of choice is provided. The chapter also explains the weather data used for the research. The chapter also discusses the information about the data manipulation process.

Chapter 6 shows the design and implementation of the dToU tariff demand response simulator. It explains the limitations and inadequacy of the dataset used for the research project. The various considerations for the simulator have been explained and the simulator algorithm and functional block diagram are provided. Finally, the output of the simulator is discussed.

Chapter 7 investigates the forecasting of demand response based on a naive model which considers a random selection of ToU tariff experiment days. Initially, a simple flat tariff energy consumption forecasting model is explained. Then the model is improvised to forecast the consumption under the ToU tariff demand response policy. The effect of the number of ToU tariff experiment days is examined. The focus of the chapter is to provide a smooth introduction to the complex problem setting of the research project.

Chapter 8 introduces the active learning methods. The ToU tariff consumption model developed in chapter 7 is used to evaluate the performance of the algorithms. The implementation of the selective sampling strategies is then shown and the results are analysed.

Chapter 9 provides conclusion of the research and provides recommendations for the further work.

Appendix A includes extra results.

2 Overview of demand response

A perfect balance of demand and supply is required for reliable operation of the grid. This means that either generation has to follow demand or vice versa. Traditionally, utility companies used to balance the grid using a vertical approach. This approach allowed the utility companies to match the demand by controlling the generation of electricity.

But later it was realised that the most efficient way of balancing the grid is by keeping the fluctuations in electricity demand minimum. Because fluctuating electricity demand results in the fluctuating generation of electricity. Changing the generation capacity of the electric grid is a complicated and costly process. Therefore, utility companies started exploring demand response strategies. Due to the deregulated markets, this new approach created various demand response opportunities which incentivised the market players to maintain the balance between demand and supply and reduce the demand fluctuations.

Demand response can be defined as 'the changes in electricity usage by end-use customers from their normal consumption patterns in response to changes in the price of electricity over time'[9]. A demand response strategy tries to lower the fluctuations in electricity demand. The electricity users are asked to change their electricity consumption pattern. To guide the users, usually, electricity retailer provides signals to their users via some communication device to guide users in required directions. These signals are often related to incentives or penalties. From the infrastructure point of view, utilities need to have a proper system in place to send the signals to their customers. For developing proper electricity price structures, they also need advanced modelling, simulation, and optimization tools, which analyze the interaction between humans, electricity infrastructure and local conditions.

With the rise of smart grid technologies [13], the coordination between users and utilities may be automatic via two-way digital communication. Demand response programs are mostly financial incentive-based energy consumption scheduling schemes which provide the users attractive opportunity to freely participate by changing their electricity consumption. From the electric grid perspective, the objective of the programs is to minimization of the energy cost of the system. Various forms of demand response have helped to maintain the reliable operation and good health of the electric grids[13].

2.1. Electricity markets

The electricity markets have both wholesale and retail components. Like any other market, in the wholesale market electricity is traded between electricity utilities and energy traders before it is eventually sold to the end customers. In the retail market, the electricity is sold to the customers by the energy retailers. The wholesale and retail markets could be traditionally regulated or competitive markets.

The traditionally regulated wholesale market means that vertically-integrated utilities are responsible for the entire flow of electricity to consumers. Usually, the generation, transmission and distribution systems are owned by the utilities in the traditionally regulated markets.

The competitive markets, on the other hand, are run by independent system operators. In the competitive markets, usually, the responsibility of the utilities is to provide the retail electricity services to customers. These utilities are less likely to participate in the generation and transmission services.

2.2. Retail electricity market

The retail markets are defined by a government body for electricity. In the traditionally regulated markets, customers do not have a choice of selecting their energy retailer or energy generator. In these markets, the renewable power plants are owned by the utilities; Therefore, it is difficult and challenging to develop large renewable power plants in the traditionally regulated markets.

In the competitive retail electricity markets, customers are allowed to choose between the retail suppliers. The degree of deregulation is subjected to vary among the competitive markets. Some markets may only allow the competitive markets at the retail end while having a state control over the power generation and transmission infrastructure. Some markets may open the power generation sector for independent producers while having control over the transmission phase. Other markets may open the transmission along with the power generation and retailing. In some markets, customers also have a choice over the selection of their power producer. For example, users can choose to get all of their energy from renewable energy providers. The competitive markets provide greater flexibility in the retail supply contract, type of generation plants and pricing of the electricity.

It is to be noted that markets may not always be divided clearly between the traditionally regulated market and competitive markets. Some markets choose to have a structure of partially competitive markets, where, only certain types of customers (usually, industrial customers) are allowed to engage in the competitive retail market. All the other sectors of the energy retailing operate as traditionally regulated market sectors.

2.3. Residential electricity market

The residential electricity market usually comes under the retail electricity market due to its distributed nature of a large number of users with small consumption. These users are usually served by an electricity retailer. These electricity retailers buy electricity from wholesale markets and sell it in the retail markets. These markets can be traditionally regulated or restructured for competitive retailers.

Conventionally, the residential electricity markets were regulated, where, the electricity is provided by a sole, usually, a state-sponsored retailer. In such a case, users do not have a choice of selecting the electricity retailer. This type of environment does not usually provide a competitive environment necessary for innovation-based growth. The traditionally regulated markets for residential users may not be driven by profit maximization strategy.

On the other hand, when residential users are given a choice of choosing their electricity retailer in a competitive market, the competition supports innovation and newer technologies which improve the efficiency of distribution.

Next section will explain the residential user demand curve and the problem of peak power generation.

2.4. Residential energy consumption pattern

Many of the residential electricity markets offer the electricity under a fixed tariff. An electricity tariff is the per unit (KW-hr) cost of electricity. Which means, that users will be charged equally throughout the day, irrespective of the time of use. This tariff policy provides an advantage of simple and clear calculation of energy consumption cost for an average user.

Although this method helps increase the transparency in the billing process of electricity usage, it keeps the major problem of electric grid unsolved; Mostly, the electricity generation plants jointly to generate the electricity to match the demand. For an electricity grid, demand and supply should always be matched. This means that, during the peak electricity demand period, the electricity generators need to generate a higher amount of power to match the demand. Figure 2.1 shows a typical residential user demand. Typically, residential customers have peak electricity consumption during the evenings. As we can see these peak demands typically last for a short time (usually 3-5 hours a day).

The power generation plants have varied capacity and response time. Response time is a time required by a power plant to match the new level of power generation. This response time differs from few minutes to few hours for various power plants. The response rate of a power plant decides the flexibility of a power plant at the time of serving a fluctuating demand. For example, if the response time of a power plant is slower than the change in the electricity demand, it creates imbalance situation on the electricity grid, risking the grid stability.

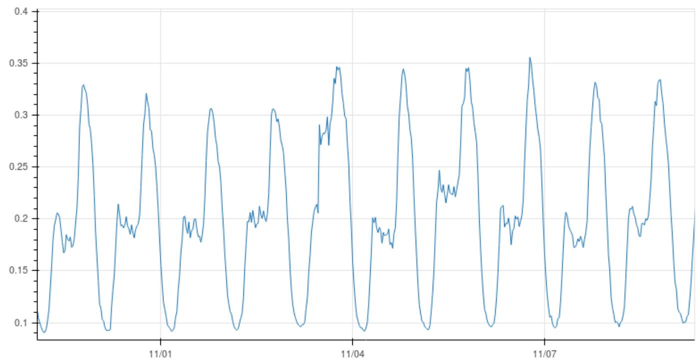

Figure 2.1: Usual residential user consumption pattern

In such cases, a spinning reserve is used. A spinning reserve is a generating capacity available to the system operator within a short interval of time to meet demand in case a generator goes down or there is another disruption to the supply. Figure 2.2 shows the generation distribution across the energy resources for the UK. We see that the base electricity generation is handled by nuclear power plants which have a slow response rate. Whereas, the peak demand is mainly handled by the gas power plants which has a quicker response rate. But, the cost of operation of power generation is high for gas-based power plants; which increases the cost of electricity for such plants. For a peak demand, the cost of generation of electricity is the weighted sum of the cost of operation (and fuel) for all the power sources. This makes the electricity at peak demand very expensive. Peaks of electricity demand are usually undesirable as they increase the stress on the electricity grid infrastructure increasing the cost of infrastructure and operations.

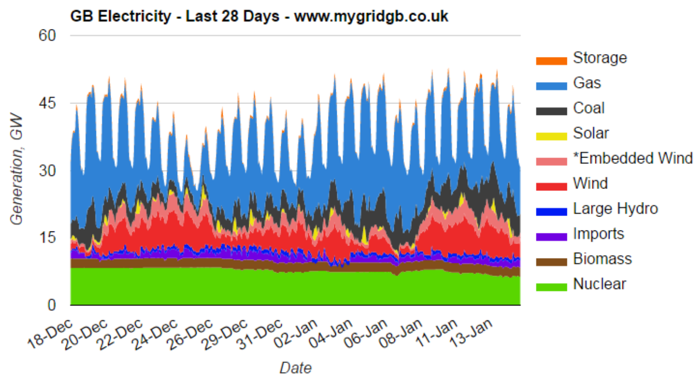

Figure 2.2: A 28 days electricity generation to match the user consumption, sorted by power sources[1]

This higher cost of electricity is directly or indirectly transferred to the end-users. It is directly transferred to the end-users when the end-user buys electricity in a wholesale market. On the other hand, if a user buys electricity from an energy retailer, the retailer sets the prices of electricity such that the actual cost of electricity is compensated in their tariff policy (fixed or dynamic).

2.5. Demand response strategies

The following section provides details about various types of demand response strategies.

2.5.1. Control mechanism

The following section discusses the classification of the market based on the control strategy of demand response[25].

Centralised control

Most of the traditional demand response schemes come under centralized control. The centralized control method consists of a central operator (usually utilities) who collects information about the customer response, electricity grid condition and then create a demand response scheme for the customers. Usually, the centralized control of demand response shows a slower response. This control can not follow fast changes in the electricity grid conditions due to the hierarchical nature of the control. But the predictability of the centralised control of demand response is much higher than decentralised control.

Decentralised control

In decentralised control, the demand response scheme is not centrally designed. In the decentralised control, the function of utility or grid operator is to just provide the information about electricity prices. The customers of decentralized demand response strategy only receive the electricity prices from the operator. The prices are a function of aggregate user demand. users change their consumption pattern based on the prices. The demand response in the decentralised case is much quicker than the centralised control. But the response of the customers for decentralised control is complicated and predictability of the decentralised control is lower than that of centralised control.

2.5.2. Demand response schemes

Demand response schemes [48] [3] are classified into incentive-based programs and price-based programs. The incentive-based programs are more useful for industrial users while price-based programs are more useful for residential users.

Incentive based programs(IBP)

In classical IBP, the participating users get incentives or discounts for participating in the demand response schemes. In the direct control, utilities can control and manage the consumption of the participating loads during the peak demand time blocks under short notice. The examples of such loads are air conditioning, heating equipment etc. These loads can be controlled by the utility to maintain the balance on the grid. As with Direct Load Control programs, customers participating in Interruptible/Curtailable Programs participants are paid with a fixed amount. The participants are then asked to change their load to a required value. The participants who do not follow the instructions are penalized under this scheme.

In market-based programs, participants do not receive fixed payments, but receive payments for their performance. In demand bidding program the customers have to bid in the market for a specific amount of load reduction. The customers have to shave their loads according to their bids. If they do not follow according to the accepted load reduction request, they are penalised. If they follow the request then they are paid for their performance. In emergency DR, the users get paid for changing their loads during the emergency situation. This scheme of demand response is beneficial for the users who have a dedicated flexible load ready to be operated. Users need to react with a fast response as the signals for such demand response scheme is provided on very short notice. Capacity market schemes ask customers to reduce the consumption only in the case of grid congestion. These programs incentive users when a part of the grid is facing contingency issues. Ancillary services provide incentives to users who change their loads based on the spot prices.

Price based programs (PBP)

The price based programs (PBP) offer dynamic price electricity for various grid conditions. These programs are as follows - the Time of Use (ToU) rate, Critical Peak Pricing (CPP), Peak Time Rebates(PTR), and Real-Time Pricing (RTP).

The Time-of-Use is one of the most popular and simple type of PBP. In the ToU scheme, the customers are provided a dynamic tariff policy to affect their consumption pattern. The ToU tariff policy usually consists of a two or more price levels which are designed such that they will follow the real time prices of electric grid. For example, the during peak consumption period, the users are provided with high tariffs. This type of price based motivation helps utilities to change user behaviour as per the requirement. CPP is similar to ToU tariff policy but it is only implemented few times a year for peak demand period. Usually the price rates of this scheme are higher that highest ToU tariff prices as these signals are only provided at critical conditions of the grid. Peak Time Rebates(PTR) asks users to shave the loads for high tariff prices on particular days. These prices are valid for entire day. In the real time pricing (RTP) scheme, customers consume electricity based on real time prices. This scheme is very similar to the decentralized control scheme shown above (only the purpose of the classification differs in both the cases).

2.5.3. Customer side strategies

The electricity user response can be achieved by three major actions[3]. Each of the response is a combination of the signals sent by the utilities and available flexible loads of users.

Load shaving

Firstly, the customers can reduce energy consumption through load shaving strategies, i.e. during the peak electricity demand period, users can choose to lower their loads. In this case, the customers do not need to change their consumption behaviour during other periods (off-peak periods). This type of response involves the temporary loss of comfort for the participating electricity users. This response is most suited for residential and commercial users.

Load shifting

Secondly, customers can give another response by shifting their loads from the peak demand time blocks to off-peak demand time blocks. The example appliances for this performance are dishwasher, washing machines etc.

On-site generation

The third type of customer response is by using onsite generation. Usually, the financial benefit from such response is lower than the above two responses. But this type of response is usually implemented where the users suffer a lot by a reduction in their electricity consumption. The example of these type of consumers are hospitals, industrial users etc.

2.6. Benefits of DR

The benefits of demand response are shown in figure 2.4. T

Economic benefits

The participating users can benefit from the demand response policies with incentives and reduction bill payments. Along with peak load prices, many dynamic tariff policies allow users to use their loads to different timings for lower tariff price. This means that users don't have to reduce their load consumption but just shift them to an alternate time period.

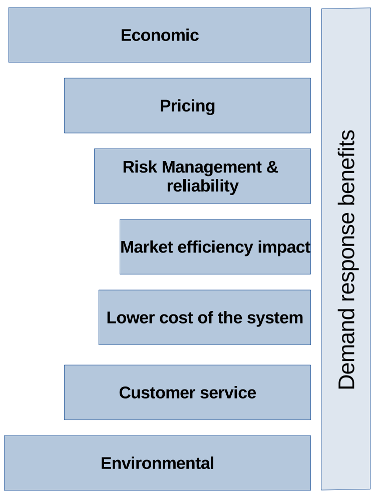

Figure 2.3: Benefits of demand response (DR)

Risk management and reliability

The demand response strategies improve the reliability of the grid as the proper allocation of penalties and incentives reduce the likelihood of extreme grid stability.

Reduction in price volatility

Another advantage of the demand response is the reduction in price volatility. As the users are guided for the demand response, the responsiveness of the users reduces the electricity price volatility. Figure 2.3 shows the relation between the price of electricity and the quantity of electricity produced.

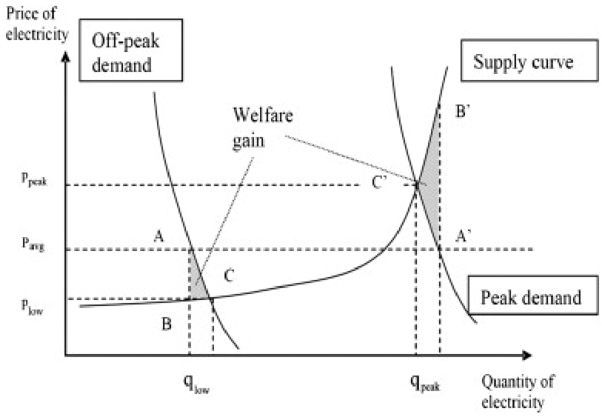

Figure 2.4: Reduction in price volatility by demand response schemes[2]

Improvement in market efficiency

The demand response programs allow the participants of the electricity grid to benefit from the actions which improve the grid stability and performance. The financial incentive motivates improving the current systems and implementing innovations which further improve the efficiency of the market.

Reduction in the cost of electricity grid

In the long run, the demand response reduces the cost of the system. As demand response schemes create transparency in the operation, it motivates to change or remove the inefficient components of the electric power system. This includes inefficient machinery, spinning reserves and other less efficient activities.

Environmental impact

The energy sector is one of the largest sources of pollution. Removing such inefficiencies reduces the wastage of energy and therefore reducing the environmental impact of the operations.

2.7. Demand forecasting

Understanding the effect of ToU tariff prices on electricity demand is a critical part of the operation of an electricity retailer. That is why the retailer has to employ forecasting models to forecast the demand for electricity for the given electricity tariff policy. The selection of these methods usually depends on the availability of the data. Depending on which data is available, the forecasting model can be based on a deterministic model or stochastic model or a combination of both. To improve the accuracy of the forecasting, usually, more than one type of models are used. The electricity forecasting works for multiple time periods - one model may forecast hour by hour consumption, whereas others may forecast the monthly or annual electricity consumption. Usually, short term models are used for the electricity trading purpose and long term models are used for forecasting the future trend and planning of the infrastructure. In this section, we will look at the commonly used methods for forecasting of electricity.

2.7.1. Trend forecasting

The trend forecasting modelling falls under non-causal modelling. In non-causal modelling, the effect of the value of variables on the predicted variable is not explained. In this method, the output target variable is expressed purely as a function of time, rather than the other factors like - economic, policy, demographic and technological variables etc. This type of trend forecasting which is obtained as a function of time is best suited for a short term forecasting application.

As the method is less complicated than the other advanced modelling techniques, the main advantage of the method is the simplicity and ease of use and understanding. However, as the function is only dependent on time, it ignores the possible effect of other variable conditions. For example, the role of weather, pricing, population size, policy change etc, is ignored by the method. The trend analysis assumes that the past pattern of a variable will continue in the future irrespective of the change of the other factors. However limited is the forecasting accuracy of this method, it provides a preliminary estimate of the forecasting of the target variable. As this method does not consider the influence of the other methods, it is a useful cross-checking tool for other short-term forecasting methods.

2.7.2. Timeseries method

A timeseries is defined to be an ordered set of data values of a certain variable. The timeseries models are the econometric models where the feature input variables used are an auto-regressive component of the timeseries. The timeseries modelling approach assumes that the future target variable will be a function of past observations of actual and predicted target demand consumption. Therefore, for making a timeseries model, we require historical data of electricity consumption observations.

One of the most valuable applications of timeseries modelling the short-term forecasting of the electricity demand. For long-term modelling econometric models are preferred over the timeseries models as timeseries models do not consider the economic factors. Another advantage of auto-regressive timeseries models is that it is simple to construct. As these models do not require any data from multiple sources, observations of the target variable are sufficient for the modelling.

One disadvantage of the method is that it does not consider the cause and effect relationship of other variables. Therefore, auto-regressive timeseries model can not make the appropriate relationship with the external factors which affect the electricity consumption of the users.

To address this issue, a multivariate timeseries models are used. In the multivariate timeseries models, instead of just using past values of the target variable, the model also uses the other input variables such as calendar events, forecast of weather etc. This addition of new variables offers previously discussed 'cause and effect' relation for the target variable.

The following section discusses the demand forecasting case study of a recent paper[6], which attempts to match the target electricity consumption by providing ToU tariff signals to a specific fraction of users who are randomly selected.

2.8. Case study: ToU tariff user selection policy

Electricity management is classically performed by forecasting the demand and adjusting the production accordingly. Smart meters or connected meters have created new opportunities for electricity management. Smart meters improved the data quality for the electricity grid, improving the forecasting of loads [10]. This has allowed non-intrusive load monitoring for industrial, commercial and residential cases [52]. But demand response for a residential case is still an interesting domain as demand response is still based on human preferences. Some preferences might be logical and can be identified by simple cause and effect analysis of common variables e.g. higher thermal loads during winters. Other preferences may not be so easy to identify and therefore create a potential risk of unknown response[28]. Therefore, performing the dynamic tariff experiments is necessary to provide information about the causal relationships with the other variables which describe the local conditions.

Target tracking method with contextual bandits approach is shown in [6]. They implement contextual bandit approach for incentive-based demand-side management. A target is set for each round and the mean consumption is modelled as a complex function of the distribution of tariff prices and other contextual variables like weather, temperature and so on. The paper tried to predict the fraction of users required to match the target consumption. The mean consumption observed is equal to

$$Y_{t,p_t} = \sum_{j=1}^{K} p_{t,j}\phi(x_t, j) + \text{noise}$$

where, ϕ is some function associating with a context x_t and a tariff j providing expected consumption $\phi(x_t, j)$. At instance t the electricity provider sends tariff j to a $p_{t,j}$ fraction of the costumers. The paper considered the population to be homogeneous, and therefore, it is unimportant to know which customers receive a tariff signal.

The above model approaches the demand response with a contextual bandit method, where they assumed that a target consumption pattern will be available for the tracking purpose.

2.9. Summary

This chapter provided an introduction to demand response strategies. The various elements of the residential electricity distribution are discussed. The various demand response strategies and their description are provided. Then the benefits of demand response are discussed with the focus on residential users. Then various forecasting models and their properties are discussed. The advantages and disadvantages of the models are provided. we discussed multivariate timeseries modelling as one of the data-driven approaches. Then we have seen a case study of a demand forecasting approach presented in [6] and discussed the assumptions made in the target tracking model.

This research project approaches the forecasting problem from a different perspective. In this research project, the number of participating users is kept constant. Instead of tracking the target, the author is more interested in improving the forecasting model by performing the ToU experiments on 'more informative' days. The selective sampling methods for selecting 'informative' data points are discussed in the next chapter along with the machine learning basics.

3

Overview of machine learning

3.1. Machine learning

In the last two decades, machine learning has become a buzzword in the computer science community. We can trace the academic interest in machine learning and artificial intelligence to the 1950s. The trend in overall growth in AI-related interest, in terms of academic publications, is shown in figure 3.1. The figure shows the number of publications which contained the keywords related to AI such as artificial intelligence, machine learning, deep learning, k-means, clustering and SVM. We see exponential growth in the interest starting from the 1990s in AI-related research. The rise of the internet along with exponentially decreasing prices of semiconductor devices have created enormous amounts of digital data and also allowed the use of more and more powerful computing at an affordable cost.

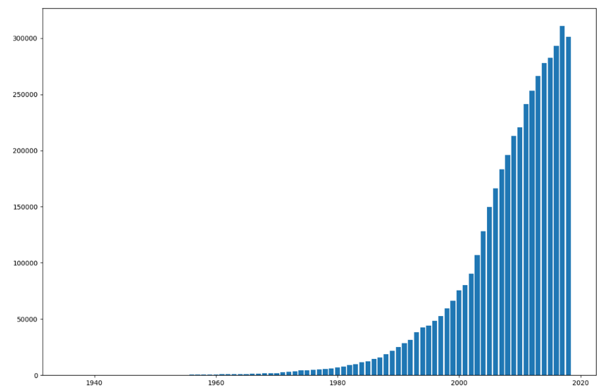

Figure 3.1: Number of publications per year containing following keywords : artificial intelligence, machine learning, deep learning, k-means, clustering and SVM.

This section will provide a basic overview of machine learning methods and terminologies for a novice reader. Machine learning is a field of study which studies algorithms and various statistical methods to perform specific tasks without any explicit instructions. The term machine learning was tossed by Arthur Samuel in 1959 [33]. There is a significant overlap between statistics and traditional machine learning methods. The more precise definition of the term is given by Mitchell [27] by the specific way machine-based 'learning' takes place: "A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P if its performance at tasks in T, as measured by P, improves with experience E.".

Machine learning algorithms are broadly classified into classification, clustering and regression algorithms. All of these algorithms support 'supervised learning'. Supervised learning algorithms build a mathematical model based on the sample set of inputs and desired outputs. Most of the times, it is easier to get inputs than desired outputs. Therefore, for classification, we can have a partial set of desired outputs for the given inputs, it is called as semi-supervised learning. In this method, some portion of the input data does not have the desired output labels. The output labels can be chosen randomly or with some selection strategy (next section covers the selection strategies). In unsupervised learning, a mathematical model is developed only based on input variables and no desired output labels. Unsupervised learning algorithms are used to find structure in the data, like grouping or clustering of instances. Unsupervised learning can discover patterns in the data. But an unsupervised learning model still has limited learning capacity as a real output of the given input is never introduced.

3.1.1. Random forest regression

Random forest is a bootstrapping algorithm with decision trees. Random forest builds multiple decision tree models with different training samples and different initial parameters. Hence, the final prediction of each decision tree is a function of these samples and model parameters. Therefore, the prediction of each decision tree differs than other decision trees of the model. The final prediction (output) of the random forest model is the mean of the predictions of the decision trees. In general, the more trees in the forest the more robust the forest becomes. In the same way in the random forest regression, the higher the number of trees in the forest gives high accuracy results.

Decision tree learning

The decision tree is also considered as a rule-based system. The decision tree algorithm will create some set of rules for the provided training dataset of target variables and features. Then the same set of rules are then used for predicting the target variables from the test dataset. The information gain and gini index are calculated [30] to calculate the nodes and formation of rules. A random forest algorithm is built on the slightly varied concepts of decision trees. In the random forest instead of calculating information gain or gini index for calculation of roots nodes and splitting of the features, the process randomly takes place.

Random forest is a combination or ensemble of tree predictors where, where each tree is a weak predictor which is based on the independently sampled random vector of features. These weak learners create an ensemble to form a strong learner. An important property of a random forest is that they do not overfit.

Decision trees tend to overfit the data by growing deep. In other words, they overfit the training data i.e. each tree forecaster has low bias and high variance. The random forest includes multiple numbers of deep decision trees, which are trained on a different part of the training data. The final prediction value of the random forest is an average of multiple trees. The random selection of training data for training of different trees reduces the variance of overall random forest model [11]. This method has a disadvantage of reducing bias and loss of interpretability, but the method usually boosts the performance of the final model.

Bagging

Bootstrap Aggregation (or Bagging for short), is a simple and very powerful ensemble method. An ensemble method is a technique that combines the predictions from multiple machine learning algorithms together to make more accurate predictions than any single considered machine learning model. Bagging can be used to reduce the variance for those algorithms that have high variance. One of such algorithms that have high variance is decision trees. An individual decision tree usually overfits the data on which it is trained. This means that for a different set of training data, the resulting decision tree will be very different, providing completely different predictions. The bagging is a method which is used to improve the accuracy of an ensemble of high variance machine learning models.

If we use a random forest regression model for prediction of a continuous type of target variable, the bagging method will consider a mean of the output of the regression trees as the final output of the model. Whereas for classification, it will consider most frequently predicted class. Bagging can also be used for classification and regression. Random Forests are an improvement over bagged decision trees.

Simple decision trees perform greedily. Decision trees try to minimize the error at every split point to select the variable. Unlike decision trees, random forest model does random sampling of features instead of selecting the most optimal split-point for the feature variables.

The algorithm

The random forests algorithm (for both classification and regression) is as follows [23]:

1. Draw n_{tree} bootstrap samples from the original data

2. Grow an un-pruned for every input sample such that at each node, rather than selecting the best split among all the predictors, randomly select a subset of predictors and choose the best split from among those variables.

3. Predict the test data by aggregating the results of n number of trees. In the case of classification, majority voting is considered. While, in case of a regression problem, the mean of all the prediction is considered for a forecast of a continuous variable.

Variable importance

Once the random forest algorithm is trained, we can calculate the drop in error function (error at a split and output) for a variable at every split point. Importance for each variable is calculated by averaging the drop in error functions for those variables. The average drop in the error at each split point and output are directly proportional to the variable importance index. For a regression forest, the error function is calculated by taking sum squared error and for a classification setting it will be gini score[23]. This procedure helps researchers identify the subset of input variables which maybe most or least important for the prediction model. The selection of important features can be performed from the above knowledge.

Another important machine learning model considered for the implementation of this research project is one class SVM classifier. The next section provides the theoretical background of SVM.

3.1.2. Support Vector Machine

Support Vector Machine or SVM can create a non-linear decision boundary using the projection of the data through a non-linear function ϕ to a space with higher dimensions. By this operation, the data which could not be separated in their original space I, can be projected to a feature space F, where, the data can be separated by a hyperplane which can separate the data points of one plane from another. When the hyperplane is projected back to the input space I, this hyperplane forms a non-linear curve. This hyperplane is represented by following equation[42],

$$w^T x + b = 0$$

where, $w \in F$ and $b \in R$. The margin between the classes is determined by the constructed hyperplane; the data points for one class are on one side, and all the data points for another class are on the other side of the hyperplane. The distance of closest points from the hyperplane for both the classes is equal; That is why SVM is also known as a margin maximization algorithm.

Slack variables ξ_i are introduced to allow some data points to lie within the margin; this reduces the chances of SVM getting overfitted by classifying the noisy data. The constant $C > 0$ is defined to determine the trade-off between the maximization of margin and number of training data in the margin (soft margin) The following minimization formulation shows the objective function of the SVM classifier[42]:

$$\min_{w, b, \xi_i} \frac{||w||^2}{2} + C \sum_{i=1}^{n} \xi_i$$

For the one class classification algorithm, the quadratic programming minimization function is slightly different than the above representation [35],

$$\min_{w, b, \xi_i} \frac{||w||^2}{2} + \frac{1}{\nu n} \sum_{i=1}^{n} \xi_i - \rho$$

Subject to:

$$(w.\phi(x_i)) \geq \rho - \xi_i$$
$$\xi_i \geq 0$$

for all $i = 1,...,n$. In the previous formulation the parameter C decided the smoothness. In this formula ν is the parameter that characterizes the solution as,

- ν decides an upper bound on the fraction of outliers from the training data (also known as out-of-class data points)
- ν creates a lower bound on number of training samples of the SVM

When Lagrange's multiplier is used to solve this quadratic problem, the decision function (classification) rule for a data point x then becomes,

$$f(x) = \text{sgn}(\sum_{i=1}^{n} \alpha_i y_i K(x, x_i) + b)$$

For one class classification, by using Lagrange techniques and using a kernel function for the dot-product calculations, the decision function becomes [35],

$$\begin{aligned} f(x) &= \text{sgn}((w.\phi(x_i)) - \rho) \\ &= \text{sgn}(\sum_{i=1}^{n} \alpha_i K(x, x_i) + b) \end{aligned}$$

The SVM algorithm is a powerful model which can separate non-linear data. With the above approach of SVM one-class classifier, we will attempt to find the out-of-class data points from our data.

But how do we optimise the machine learning model's performance? We can optimise it either by fine-tuning the model parameters or optimising the data. Optimising the model parameters has gained high momentum in the early days of machine learning, but in recent years, there is a growing interest in optimizing the data for improving model performance. Active learning is one of such methods. The following section explains how active learning, a special case of machine learning selectively samples the desired output labels.

3.2. Introduction to active learning

Active learning is a special case of semi-supervised learning where the instances for desired output label are selected with some 'intelligent' strategy. One of the most popular 'intelligent' strategies is to evaluate the importance of data based on how much information it contains. For the following discussion, An 'instance' is defined as a data point received by an algorithm, whose informativeness has to be measured. The term 'query' is defined as a request created by an algorithm for labelling of an instance. The algorithm queries desired labels for a limited number of inputs to reduce the cost of data. But what is meant by 'cost' associated with data? Let us go through an example. In the medical field, we can use machine learning algorithms to classify MRI (Magnetic Resonance Imaging) images as positive or negative for a particular disease. To train a machine learning model, we need input data (i.e. MRI image) and desired output data ('positive' or 'negative' tag). In this case, getting thousands of MRI images is quite cheaper when compared to getting labels for those images from professional medical doctors. Therefore, in such cases, if we want to optimize the cost of the total sample set, we need to reduce the number of queries we send to the doctors to label the output based on input tags. Mind that the problem we are solving is not focused on cost reduction but constrained by it. The goal of the active learning model is to achieve maximum accuracy with a limited number of output labels. So, if we need to maximize the total accuracy of the model, we need to create a 'curious' machine learning model which intelligently selects the instances whose labels will maximize the information gain. We can create a specific strategy to choose such instances based on certain characteristics of machine learning models and information theory. The study of these strategies is called active learning.

For many supervised learning tasks, getting a label is an expensive, time consuming and/or difficult process. Active learning tries to overcome this labelling bottleneck by asking queries for selected unlabeled instances to be labelled by an oracle (i.e. a human expert). This way, an active learning algorithm tries to improve the accuracy with a small number of labelled instances, hence reducing the total cost of obtaining the labelled data.

But how does the current established problem of this thesis would be benefited from the active learning techniques? As we discussed in previous chapters, the issue with deploying wide-scale dynamic Time of Day tariff is that we do not know the effects of such implementation. So first, we need to design some trial phases to understand user response to such tariff experiments. Setting up such trials is no easy process. It involves the distribution of costly smart meters, design and implementation of response measurement setup, extremely high coordination between various utility entities and most importantly, unknown impact on the interconnected national/international grid. It will be an understatement to call these trials anything less than elaborate multi-player scheme of complex experimental setups. There is a high cost associated with every action performed under such trials. Therefore, one of the constraints of such trials is to optimise the cost of the experiments. We can use aforementioned active learning techniques to minimize this cost by reducing the number of total experiments while improving the accuracy of the machine learning model which learns the demand response of tariff signals under given conditions of the input for any day. By using a machine learning model for forecasting, we can further reduce the dependence on the live trials for understanding future demand response.

For the selective sampling of an experiment day, we can use various methods explained in section 3.3. One way of approaching the selective sampling problem is by a concept called 'informativeness measure'. The idea is to calculate a measure which indicates the 'information entropy' of an instance. The instances with more informativeness are often queried for labelling. Information entropy is a function of negative logarithm of probability mass function of the value: $S = -\sum P_i \log P_i$. So when the probability of an instance is low, meaning the occurrence of such instance is low, then the negative logarithm of probability function produces higher value when compared with high probability instance. information entropy refers to uncertainty or disorder, and the concept is widely used in information theory originally proposed by Shannon in 1948 [39]. Although, this is not the only method of quantifying the 'informativeness' of an instance. Following chapter will provide other means of calculation of 'informativeness'.

Now as we know the role of active learning in the current project, let us dive into active learning theory. In the active learning, we have three major scenarios where a learner (machine learning model) can ask a query to label instances to an oracle (can be human or machine). These settings are given as - i) Membership Query Synthesis[4], ii) Stream-based selective sampling[7] and iii) pool based sampling[8]. Figure 3.2 shows the illustrative framework for the above-mentioned settings.

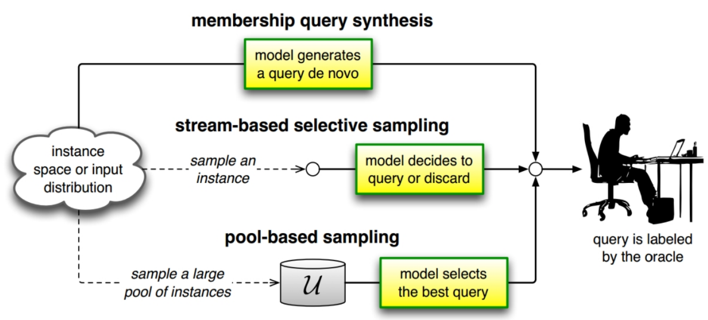

Figure 3.2: Illustration of the active learning settings by Settles [36]

3.2.1. Membership query synthesis

We will start with a membership query synthesis framework discussion. Membership query returns only the following information: if the given input instance is part of an unknown set or not. One of the first elaborate discussion on the topic is traced back to 1988 [4]. The model can query a label for data point from oracle. The query data point is synthetically generated by the learner. For a non-uniform distribution, this creates a problem. As the algorithm synthesises a query based on its hypothesis, the generated query may not follow the underlying distribution of data. This will result in the generation of an instance which may not make sense in the real world. For example, let us consider the classification problem of handwritten numbers. If a membership synthesis learner queries a synthetic image, the resulting 'digit' of the image may not make any sense to an oracle. This results in the generation of garbage data.

Jackson [17] shows the efficient and successful implementation of membership query function for disjunctive normal form (DNF) formulas. Whereas, Valiant [47] provided a negative result of the membership query method and showed that general Boolean circuits are unpredictable under membership query synthesis. He stated [47] that the learner's knowledge about unnatural inputs is irrelevant and there is a requirement for the learner to know every possible combination of the inputs.

But we can infer from the above discussion that this framework is not suitable for the proposed project as it can only work infinite problem domain. For example, if we decide to use the above framework on our problem, the learner will generate a synthetic instance (energy consumption for a day) for which it has less clarity of the concept. We have no way to find out the demand response for such a hypothetical day to improve the concept of learning. As it depends on the generation of synthetic data, we can not incorporate the framework in a real-time system, where the inputs can not be controlled.

3.2.2. Stream-based Selective sampling

The second type of active learning framework is stream-based selective sampling, where an instance is selected for labelling, literally from a stream of samples. The learner one by one decides either to accept or discard the sample instance for the labelling process. Usually, a collection of unlabelled data is inexpensive. But labelling the data incurs high cost. Therefore, stream-based active learning is a useful framework for a scenario where the learner can have a continuous stream of data. The method analyses every instance of the stream. If the algorithm is uncertain about an input instance, it queries label from an oracle.

For the uniform distribution of input samples, the stream-based selective sampling performs similar to membership query learning. But the real advantage of stream-based selective sampling is that for a non-uniform and unknown distributions query will be sensible as they are drawn from a real underlying distribution.

Various approaches to form a query are discussed section 3.3. The information entropy-based implementation is seen in Cohn [7]. Cohn handled the idea of selective sampling with a neural network, where for binary classification, a regression model was used which provides output between [0,1]. Here 0 represents one class and 1 represents another class. By allowing continuous values between 0 and 1, the uncertainty was quantified. When the model predicts value which is closer to either extremity, the model shows high confidence on the concept that it has learned from training data. Whereas, when the output lies near 0.5 value, the model is uncertain about either outcome. Therefore, from an active learning perspective, these instances make a curious case for the exploration. The query generation, in this case, can be handled with multiple criteria. The simplest criteria are to compare the model output with a threshold.

Stream-based selective sampling can be approached by multiple methods. Zhu [50] showed the active learning mechanism for a stream of data which is usually stored in chunks. The implementation involved labelling of instances from a batch of data streams using uncertainty sampling. This kind of implementation assumes that limited re-access to the data is possible. Another way of stream-based query selection is performed in an online setting, where the labelling decision is made at the time of scanning of each instance. The later type of implementation is found in Žliobait [51]. They also handled the problem of concept drift using uncertainty sampling and dynamic allocation of labelling efforts by randomization of search space.

Thompson used active learning for natural language parsing and information extraction [44]. They claimed to achieve upto 44% annotation cost for speech dataset. [40] used stream-based active learning for sentiment analysis in stock market prediction using a stream of twitter data.

In the last few years, interest in pool-based sampling has been renewed as computer vision and image-based artificial intelligence has seen tremendous growth. The pool-based scenario is explained here.

3.2.3. Pool-based sampling

Lewis and Gabe [22] introduced pool-based selective sampling method in 1994. Today, it is very common to have a large pool of unlabelled data. A large collection of digital images is a good example of such a pool. This type of domains inspires the pool-based methodology. In this method, a learner which has access to a pool of data can query labels for a limited number of unlabelled samples. This type of active learning is more popular than the other two types shown above as nowadays it is possible to have a large quantity of unlabelled data readily available.

In a general pool based method, the active learning algorithm is used to calculate the informativeness of all the instances. There are several ways of calculating informativeness of an instance, which will be discussed in section 3.3. Then, the queries are selected in a greedy manner where the instances with a higher value of informativeness are selected first. The method can be implemented iteratively.

A pool-based recommender system is proposed by Rubens [32]. Tong used support vector machine (SVM) based active learning for text classification [46]. Joshi [18] has shown a pool based sampling method for image classification using SVM based probability estimation of class membership. The pool-based selection method finds a lot of applications in the multimedia domain as the arrangement of a large pool of data is possible for most of the multimedia applications. A literature survey by Wang [49] focuses on the active learning techniques in the multimedia domain. A classic literature review by Settles [37] shows a large number of examples of the pool-based method in the application domain of images, text, audio, speech, video etc.

Out of the frameworks discussed above, stream-based and pool-based methods query a real instance for labelling from the distribution, whereas membership query function will only query from real distribution if the distribution is uniform. The main difference between stream-based and pool-based active learning is about the availability of the data during query synthesis. In stream-based learning, data is sequentially scanned and query decision is made individually. Whereas, in pool-based learning, the complete set of available data is used to rank the complete dataset for selecting the best query. The pool-based approach is much useful when working with offline dataset or availability of large computation power and large scale storage (e.g. servers and cloud computing), whereas stream-based method works best for online situation (e.g. ToU tariff response problem of this thesis) or limited power and storage capabilities (e.g. mobile and embedded devices).

3.3. Query strategy framework

All of the active learning frameworks evaluate the informativeness of an unlabelled instance. This section provides an overview of the general strategies from the literature. From now on, x_A^* is used to refer to the most informative instance according to some query selection algorithm A. In the scope of this project, the following are the related methods for obtaining the informativeness of input instance.

3.3.1. Uncertainty sampling

The label sampling under this query strategy is done by selecting the samples for which the model is least certain about. This framework decides the informativeness of an instance based on the certainty with which algorithm can label the instance. The approach is quite straightforward for a probabilistic model for binary classification. For example, in a binary classification, for a particular instance, if a probabilistic model shows 0.5 probability for both the classes, then the model is completely uncertain about the label of that instance. The assumption is that knowledge about labels of these uncertain instances will potentially improve the decision boundary of a classification problem.

Therefore, the concept of uncertainty sampling revolves around finding the instances which are closer to the decision boundary, therefore harder to classify. In the SVM based classifier, Tong [46] proposed a strategy which queries the instance which is closest to the decision boundary of SVM classifier.

Another popular uncertainty sampling measure uses information entropy [39] as an uncertainty measure:

$$x_H^* = \arg\max_x - \sum_i P_\theta(y_i|x) \log P_\theta(y_i|x)$$

where y_i ranges over all possible labels. Higher the amount of information entropy more is the information in the distribution. In machine learning, this concept is used for calculation of uncertainty which is directly related to the informativeness of an instance. The entropy is equivalent to a margin approach for a binary classification problem.

The uncertainty sampling methods for three or more classes are shown in Settles [37]. As the scope of this thesis only has requirements for binary classification, we will limit the study of uncertainty sampling for binary classification problem.

Let us consider the issues related to the uncertainty sampling approach. As this approach always searches instances for which the model is least certain, it falls into the trap of greedy search, limiting the exploration. Huang [15] found that the margin-based methods (as shown in this section) do not perform well at the beginning of training session as for the lower number of training samples, the decision boundary is often inaccurate.

3.3.2. Query by committee

Query by committee approach considers the ensemble of models for deciding query of a label for unlabelled data. Let us consider a set of n sub-models - $f_1(x), f_2(x), ..., f_n(x)$. Where, x is the input data and $f_i(x)$ is the output function of an ith sub-model. The ensemble of models is usually trained on a slight variation of training data. The criteria for selection of the training data for each model in ensemble depends on a method to method.

As every sub-model of the above ensemble is trained on slightly different data, each model has a slightly different hypothesis. The hypothesis of a machine learning model is built on the training data. This results in some degree of disagreement between the sub-models. The Query by committee method tries to quantify the disagreement between the members (sub-models) of the committee (ensemble).

The QBC algorithm achieves the disagreement between the committee members by minimizing the version space. The concept of version space was proposed by Mitchell in [27]. For a binary classification problem, the committee of the n models can create n number of hypothesis. If we say that some of the hypotheses provide more 'tight' boundary for positive training samples than other hypotheses, we call such boundaries as 'specific boundaries' (SB). The other hypotheses which create boundaries which contain positive samples as well as remaining feature space without including any negative samples, we call them 'general boundaries' (GB). The goal of the query by committee algorithm is to constrain the size of the area between two boundaries, such that the search of the queries is more precise.

In 1992, Seung [38] proposed this algorithm. The paper considered a committee of two models for the classification problem. The query is chosen according to the principle of maximal disagreement [38]. The paper showed that the generalised decreases exponentially as the number of queries go infinite. Seung found that the active sampling by QBC approach showed better results than random sampling.

The method is usually implemented for binary classification. The method also assumes that the data is separable. The Gibbs training procedure is used to train the QBC model. The assumption that Gibbs training procedure will be used makes the algorithm less useful for a deterministic component learning problem. To solve this problem, [26] proposed two algorithms - Query by bagging and query by boosting to construct a committee. The results of [26] show that the query by boosting achieves best results followed by query by bagging and then traditional query by the committee for a deterministic component learning problem.

The study by McCallum and Nigam [29] showed that a small number of labels queried by active learning algorithm can reduce the prediction error by around 33%. The paper also provides a proof for the value of unlabelled data in the active learning domain.

3.3.3. Variance reduction

As the problem setting of this research project deals with the regression model, we now consider a regression setting. We can tweak the query by committee algorithm to incorporate the uncertainty of the regression instance.

Let us consider a set of n sub-models of an ensemble - $f_1(x), f_2(x), ..., f_n(x)$. Where, x is the input data and $f_i(x)$ is the output function of an ith sub-model. In this case, the sub-models of ensemble are regression models. We train each sub-model on slightly different regression data. The strategies to select the training data for these sub-models may differ model to model.

The ensemble of the above models can be used to predict the actual value of the target variable $y(x)$. The ensemble of n models is used to predict the target variable $\hat{y}(x)$ for given input x. The output of the ensemble of weighted sub-models is mathematically shown by,

$$\hat{y}(x) = \sum_{i=1}^{n} w_i f_i(x)$$

Now, as each model is trained on slightly different data, we again try to minimize the variance of the ensemble by querying the instances which provide show a high degree of ambiguity. In the variance reduction method, this ambiguity is quantified by the variance of the ensemble. Therefore, we calculate the variance V of the ensemble of n models for input x as,

$$V = \sum_{i=1}^{n} w_i (f_i(x) - \hat{y}(x))^2$$

In variance reduction method, we try to maximize the knowledge of the ensemble model by trying to minimize the variance of the model. The reduction in variance is directly proportional to the confidence of the model in certain common concept. Therefore, the objective of the variance reduction model is to minimize the generalized error by querying the label of unlabelled instances with high prediction variance. The terms 'label' and 'unlabelled instance' is borrowed from the classification domain. Under the current setting, the term 'label' means the target variable and 'unlabelled instance' means the input for which the target variable is not available.

Now, we will discuss various ways of ensuring the difference between the hypotheses of the sub-models. One way of creating such a committee is by initiating the sub-models with different initial conditions. This can be achieved by having a different set of hyperparameters for each sub-model. Another way of achieving the difference in hypotheses is by selecting different training data for each model such that models share minimum common knowledge about the feature space. The former method was observed in neural network-based model ensemble implemented by Krogh and Vedelsby [20]. They have used a committee of neural network models with different hyperparameter settings for creating a disagreement between the models. The example of the other type of method to ensure the disagreement between the users is shown by Raychaudhari and Hamey in [31]. They achieved the ambiguity between the sub-models by providing separate training data. The paper claims that the accuracy of their model is greater than [20]. Further discussion about the claim has not been found.

3.4. Discussion: Considered active learning methods

In this section, we will formulate the two active learning methods for the problem setting of the research project. First, we will try to combine the knowledge of section 3.1.1 and section 3.3.3.

As discussed in section 3.1.1, the decision trees of the random forest algorithm are considered to be 'weak learners'. They tend to have high variance and low bias in the output. This property of the random forest algorithm can be exploited to find out the uncertainty in the output of the model. Instead of averaging the output of the trees via the bagging process, we calculate the variance of the ensemble of trees.

Secondly, we will consider a novelty detection algorithm based on SVM classifier. In this algorithm, we will try to create a decision boundary which will separate novelty data points from the base distribution of the input data. Then we will calculate the distance of every instance from the decision boundary and use that as a measure of selective sampling.

3.5. Summary

This chapter introduced the concept of machine learning. Then a special case of active learning was discussed from the research project point of view. The working of a random forest and one-class SVM classifier is discussed. We observed that active learning for classification problem setting has been well developed, but they work in the regression domain is still limited. We then analysed the variance reduction method for regression problem setting in detail. Finally, we revisited the random forest method to give details about the implementation of the variance reduction method via high variance decision tree ensemble of random forest algorithm. We also reviewed the novelty detection method using SVM. In the next chapter, we will look at the design of the experiment setup of the research project.

4

Setup design

This chapter introduces the setup considerations for the research project and discusses the elements of the setup. This is a stylised problem that aims to capture the main elements of the real challenge - electricity retailer and residential users behaviours. The assumptions and considerations while modelling both elements are considered in this chapter.

4.1. Introduction

One of the goals of the research project is to create a forecasting model to predict the energy consumption of residential users under dynamic Time-of-Use (ToU) tariff policy. The electricity retailer provides the ToU tariff signals to customers. Then the customer consumption pattern is influenced by the electricity prices, usually bringing the demand down for HIGH price signals. The research project analyses the effect of ToU tariff signals on the consumption pattern of the users and finds out the ways to improve the forecasting of the demand by the means of active learning techniques explained in chapter 3.

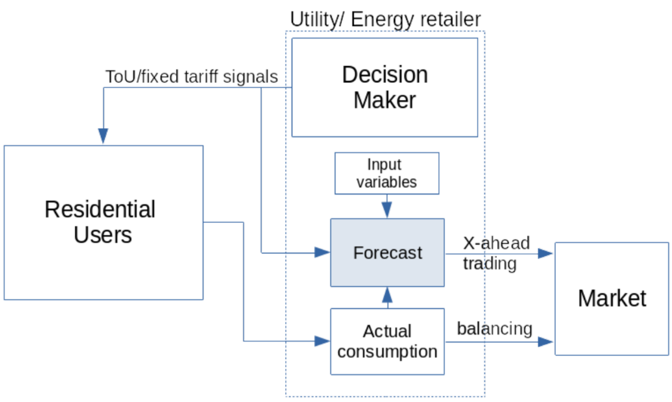

Figure 4.1: Design of the setup of the research project

Figure 4.1 shows the schematic for the design of the experimental setup of this thesis. Let us discuss each of the concepts introduced in the figure.

4.2. Electricity retailer

Electricity retailing is the final sale of electricity to the end-users. Energy retailers buy electricity from the market usually at a day-ahead price in a spot market or hour-ahead prices in an intraday market based on the forecasted demand.

Let us consider the importance of a good forecaster. The end-users expect a reliable and uninterrupted supply of electricity from their electricity retailer. As we know that the demand and supply of electricity have to match all the time. The electricity retailer bids electricity for the future time period based on the expected demand for that time period. Therefore, if a retailer observes mismatch in the actual demand and forecasted demand of the electricity, the retailer needs to ensure that the gap in the demand of the electricity is reduced using various demand response strategies explained in chapter 2.

The failure to match the ordered electricity consumption results in imbalance on the electricity grid. In that case, the responsible party has to pay the very expensive balancing cost. Therefore the market players try to avoid the imbalance conditions. In our case, the electricity retailer can reduce the balancing cost by having a good forecaster.

For this project, we are considering the case of forecasting electricity consumption under ToU tariff. The electricity retailer decides a tariff policy for a day; the policy can be a ToU tariff or fixed tariff. The expected user behaviour is that users will adapt their usual electricity consumption pattern to avoid the higher electricity charges for the peak demand time period under the ToU tariff and will continue the usual consumption under fixed tariff signals.

The peak demand prices help the utilities to reduce the demand at peak period. This is especially useful when the transmission system faces issues like constraints. Under the constraints (related to thermal, voltage or frequency) the transmission system reaches its peak power carrying capacity. In other words, the transmission systems can not cater to higher power demand if required. This condition creates instability on the grid. The utilities and grid operators try to avoid this condition. The dynamic tariff demand response strategy is one of the methods which can reduce the congestion on the grid.

Now, we will look into various elements of the electricity retailer's considered domain: forecaster, decision-maker and tariff policy.

4.2.1. Forecaster

This thesis focuses on the design and implementation of a ToU tariff consumption forecaster and decision-maker algorithm. The setting of the forecasting scenario is explained here. It is assumed that the users are never subjected to a ToU tariff policy before starting the trials. Therefore, we do not have any information about user response before the trials. The energy retailer will set ToU tariff on some days and send the corresponding signals to all the users. On other days, the energy retailer sends the fixed tariff signal to the users.

The forecasting model is then trained on the electricity consumption data gathered from the residential users with tariff signal and other variables as inputs features. The target variable is the electricity consumption of the users (under ToU and fixed tariff policies). The feature variables consist of the tariff signals and other variables which describe the local conditions. The detailed discussion on the feature variables is done in chapters 5 & 7.

Considering the discussion on the demand response forecasting models in chapter 2, we will consider a timeseries model for the forecasting purpose. As we have discussed in chapter 3, the collection of ToU tariff electricity consumption data is an expensive task, we will use a combination of fixed tariff and ToU tariff responses by the users.

4.2.2. Decision maker

The design of a decision-maker is the main aim of the research project. The function of a decision-maker is to schedule ToU tariff experiments on certain days. For the prediction of the electricity consumption under ToU tariff policy, the historical data corresponding to the same observation i.e. the observations of electricity consumption under ToU tariff policy. The decision-maker is responsible for gathering that data by performing the ToU tariff experiments.

As we are considering the cost for each ToU tariff experiment, we are constrained by the total expenditure on the ToU tariff experiments. Therefore, electricity retailer can not have an unlimited number of experiments but will have to choose specific days for ToU days. The decision-maker can use various strategies for choosing the days for the experiments. For example, in the simplest case, the decision-maker can randomly choose the days for ToU experiments without considering the impact of external factors. In some other case, the decision-maker will consider the grid conditions while selecting the experiment day. For example, the decision-maker will choose any day where the least impact on the electricity grid is expected. Still, it will be random selection under that constraint, as the decision-maker is not concerned with the actual day itself.

This research project has the following consideration for the decision-maker. All the resources required for performing a ToU tariff experiment is available to the electricity retailer when the decision-maker asks the retailer to perform the experiment.

The random selection of experiment days is considered as a passive selection criterion. If the decision-maker employs any ‘smart’ method for selection of an experiment day, then the selection method is called selective sampling or active selection criteria. The thesis focuses on finding out the active selection criteria for selecting experiment days which are more informative than others. In either case (active selection or passive selection), the decision-maker is an algorithm which fulfils the given purpose.

4.2.3. Tariff policy

For the research project, the investigation of electricity consumption under Time-of-use (ToU) tariff policy is considered. The design of ToU tariff policy is done by the electricity retailer, which includes the size of each time block, levels of the tariff, price of electricity at those levels and duration of each level of tariff for a ToU tariff day. Let us discuss each of the above-mentioned design aspects.

- Every day is divided into 48 half-hourly time blocks for the design of ToU tariff policy. The electricity consumption readings are taken every half hour (with the unit, kW-hr per half hour). The tariff policy is subjected to change at a particular time block of a day.

- The tariff prices are designed such that the residential users who do not participate in the ToU tariff response, will not be penalised.

- The considered problem has the minimum ToU tariff price levels: NORMAL and HIGH. The NORMAL price signal corresponds to a comparatively lower tariff time block, where the use of electricity is charged with a lower tariff. The HIGH price time blocks of a day correspond to the time period where the electricity retailer wants its users to reduce their electricity demand. The HIGH price is sent to all users during peak hours, whereas the NORMAL price is sent to the users at off-peak hours. Price of electricity for the NORMAL and HIGH levels is kept the same throughout the trial, meaning, the user response is not affected any time during the trials due to any change in the tariff prices.

- In theory, as the ToU tariff policy is divided into 48 time blocks, the electricity retailer can generate a tariff policy which has multiple changeovers between NORMAL and HIGH price. But by design, we will put a constraint of one HIGH price block per day. The consideration is inspired by two reasons,

  - Complex design of ToU tariff is difficult to follow for residential users, therefore, usually simple peak and off-peak time blocks are defined for a day [24]
  - Main focus of two-level ToU tariff is to reduce the peak electricity demand, which is observed only once during a day

  Due to this constraint, we will focus on reducing the peak demand of a day. Therefore, the HIGH price will be set for a day such that the peak demand period is covered by the tariff. The users will receive a NORMAL price for all the other time blocks of a day.

- The start time and duration of a HIGH price is randomly selected for a day and the resulting policy is sent to the decision-maker to analyse. If the decision-maker decides to perform the ToU tariff experiment, then for that particular day, the ‘observed’ electricity consumption for ToU tariff is collected. If the decision-maker decides to not perform the ToU tariff experiment, the generated ToU tariff is discarded and a flat tariff with all NORMAL signals is sent to the users and the corresponding electricity consumption is collected. The input features corresponding to the day under consideration is collected by the electricity retailer in either case.

This concludes the discussion on the various elements of considered electricity retailer domain. Next, we will discuss the considerations for residential users of this trial design.

4.3. Residential users

The residential users considered under this thesis are based in London. The research project is utilizing the data of the Low Carbon London (LCL) project [24]. The dataset is explained in chapter 5. The data from the Low Carbon London project is used to simulate the behaviour patterns of residential users. Therefore, the data used for the research project is the simulated behaviour of the users. The simulator design is explained in chapter 6. This section focuses on the assumptions and considerations made about the residential users,

- The residential users are a homogeneous mixture of the residents of London [24].
- The input data for the data simulator is actual observations of fixed tariff electricity consumption collected under Low Carbon London trials [24]. Therefore, the LCL fixed tariff data is used as a base data for simulating the fixed tariff electricity consumption and ToU tariff demand response for this setup.
- All the users of the given electricity retailer have selected to receive the ToU tariff signals a day before the ToU tariff day.
- It is not mandatory for the users to engage in the demand response for ToU tariff. If a user chooses not to participate, the user will continue with their usual consumption behaviour (i.e. a consumption similar to that of fixed tariff consumption)
- If the users are interested in participating (will be called as participating users) in the demand response, they will reduce their consumption during the HIGH price signal.
- Every participating user can lower their load in one step which is a fraction of their fixed tariff load. The participating users will start reducing their loads at different times to create a smoothing effect for the tariff response. Similarly, the participating users will stop responding to the tariff at different times. The delay (positive and negative) around the start and stop time of a HIGH price signal is calculated in the data simulator.
- The above participation rate is predefined for a trial. Therefore, a fraction of users will be randomly selected for following the ToU tariff, while the other users will continue their normal consumption pattern.
- The electricity retail will only receive the aggregate consumption of the users which is the sum of electricity consumption of all the customers of the electricity retailer.
- The simulated demand response of the residential users is a function of their fixed tariff electricity consumption. Therefore, for any particular day, the demand response quantity (in kW-hr) is only dependent on the value of fixed tariff electricity consumption.
- The demand response calculations will only consider peak load shaving without shifting any load. For example, when users receive a HIGH price, they will reduce their loads during the HIGH price time block and will continue with their usual consumption once the NORMAL price resumes.

4.4. Summary

This concludes the chapter 4. In this chapter, we have seen the basic setup of the thesis research work. The various aspects of the considered domain of electricity retailer are discussed while providing the assumptions of the designed setup. The implementation of the electricity retailer's forecaster and decision-maker blocks are discussed in chapters 7 & 8. The considerations regarding the modelling of residential users' consumption behaviour are discussed. The residential consumption is modelled in chapter 6. In the next chapter, we will discuss the data used for the research project.

5

Data

Data is an important part of the study of machine learning techniques. This chapter introduces the data used for the evaluation of the proposed active learning methods. The research project uses two datasets namely - Low Carbon London (LCL) ToU tariff trials dataset and MIDAS weather dataset. The information about both datasets is presented in this chapter. As both datasets partially contain redundant data, the data cleaning process is explained. Finally, the final dataset is introduced, which will be used for the research project.

5.1. Low Carbon London Project

The demand response model used in the thesis is extensively based on the public data released Low Carbon London (LCL) project, which was the UK's first residential sector, dynamic time-of-use electricity tariff trial.

5.1.1. Background

During 2011-14, UK Power Networks (the Distribution Network Operator for London) ran the Low Carbon London project which was funded by the Low Carbon Networks Fund (LCNF) run by OFGEM. This project tested many low carbon technologies on London's electric power grid [24]. Imperial College, London gathered data from all the trials which were performed during the tenure of this project and perform analysis to get insights from the data. The anonymous individual user energy consumption data was gathered using smart meters. For the scope of this project, we will only consider the residential dynamic tariff trials involved in LCL project. The following information is based on the report published after the trials[34].

5.1.2. Experimental groups

The trial participant group consisted of 5,533 household users. These users were provided with the smart meter which recorded their energy consumption during the tenure of the project. Out of 5,533 participants, 1,199 participants were chosen to participate in dynamic Time-of-Use tariff trials. The remainder acted as a control group for the trials.

5.1.3. Experimental units

The experimental unit can be considered as a scope of a single experiment. Whenever we consider an experiment which is performed on a large number of subjects, we have two main sources of uncertainty: the variation in the response between various participants and variance caused due to difference in surrounding conditions related to each day. Therefore, it is important to clarify what factors are to be considered in an experimental unit.

The idea of a trial day is to set a standard time period in which various tariff prices can be provided, independent of each other. The researchers thought that the sleep cycle of humans was a clear divider of human days. Therefore, starting a morning (05:00) was considered as the start of a new trial day

For the current thesis work, the same definition of an experiment unit is used. But to model the aggregate consumption, a bottom-up approach is used. That is based on the findings of the trials, we try to mimic the effect of individual households to obtain same aggregate level results i.e. every household response is modelled individually, so that when we consider aggregate consumption, we will see the same effect as seen in original trials (More on this topic in chapter 6).

5.1.4. Household sampling

The participants were recruited based on the opt-in process. The demographics of participants were carefully selected to match the overall demographics of London. The participants who did not further opt for dynamic tariff trials were then grouped under non-dynamic (fixed) tariff category. The fixed tariff group was a control group which provided a baseline consumption to analyse demand response of the dynamic tariff group. Figure 5.1 shows the individual user consumption of electricity under fixed tariff policy.

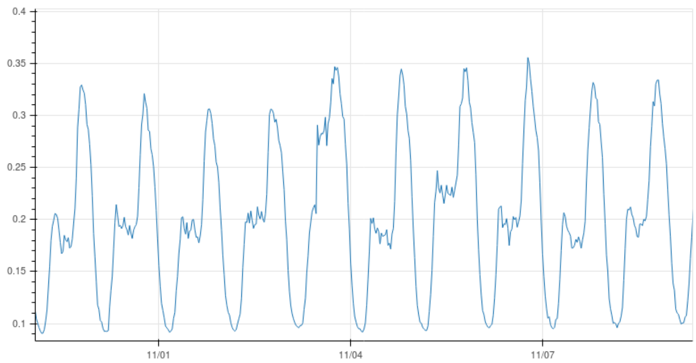

Figure 5.1: Usual residential user consumption pattern

5.1.5. Parameterization of demand response

This section describes the parameters which were used to quantitatively describe the demand response actions of users for a given dynamic tariff policy. The various parameters are given below -

- Electricity prices: The trials included following tariff price levels - LOW, DEFAULT and HIGH.
- Timing of an event: The event is selected based on time parameters such as time of day, day of the week, seasons of year etc.
- Persistent of an event: Only three consecutive trial days were allowed.

5.1.6. Overall price schedule

As the LCL trial was first of its kind trial in the UK, there was a lot of ambiguity about the effect of dynamic prices on the response. Therefore, the tariff price levels were set such that the expected user response will be maximized.

The tariff price levels are decided such that if a user decided not to participate in the trials, the net effect of their inaction will be zero. In other words, users will not be charged for not changing their consumption for the dynamic tariff. The rates for price bands were as follows:

- HIGH: 67.2 pence/kWh
- DEFAULT: 11.36 pence/kWh
- LOW: 3.99 pence/kWh

The nonToU tariff group was charged at a fixed standard rate of 14.228 pence/kWh.

The actual price levels of tariff do not affect the implementation of our research project as we assume that the rates are constant throughout the simulated trials of this research project. But it is interesting to understand the economics and incentives related to the LCL trials. We could see the effect of revenue-neutral policy on user behaviour. Around 50% users who actively opted-in, never actually participated in any trials[45].

5.2. Summary of LCL dataset

The above discussion on the LCL dataset is only limited for the scope of this research project. The following table shows the summary of the considered part of the LCL dataset gathered during the dynamic Time-of-Use tariff trials of LCL project.

Total number of users	5,533 users
ToU tariff users	1,199
Fixed tariff users	4334
Time period of the trial	One year (2013)
Tariff levels	3 (LOW, NORMAL, HIGH)
Total HIGH tariff events	21

5.3. Weather data

The weather data like temperature and weather type is an important feature for the forecasting of electricity consumption of the residential users. Therefore, we will explore the weather data as a potential set of input features.

5.3.1. MIDAS dataset

The research project uses the MIDAS UK hourly weather observation data and contains meteorological values measured on an hourly time scale. The measurements of the concrete state, wind speed and direction, cloud type and amount, visibility, and temperature were recorded by observation stations across the UK.

The MIDAS UK hourly weather observation data contains the weather data from all the weather stations from the UK. As LCL trials were performed in London, the London based weather stations are considered for the weather inputs. Due to the concerns of privacy, the exact location of the users is not provided in the LCL data. The Heathrow (London) weather station data was found most complete with some of the parameter logs lead back to 1948. Due to the reliability and availability of the data from 2011 to 2014 (the total period of LCL trials), Heathrow weather data is considered for this project.

5.3.2. Study of parameters

As they say, the machine learning task is 80% data processing and 20% fun. The UK hourly weather observation data contains 104 parameters, including timestamp, weather id, device id, weather conditions (like temperature, pressure, wind speed and dew point). Some of the important parameters are described here,

- ob_time: Date and time of observation
- wind_speed: measured in knots
- visibility: measured in decameters
- msl_pressure: Mean sea level air pressure Unit=1 hpa to the nearest 0.1 hpa. Precision aneroid barometers are now in general use for measuring pressure and a correction for altitude is applied to obtain the value at mean sea level (MSL).
- air_temperature: Unit=1 deg C to the nearest 0.1 deg C
- dewpoint: Dew point temperature is the temperature to which the air must be cooled to produce saturation concerning water at its existing atmospheric pressure and humidity. Unit=1 deg C to the nearest 0.1 deg C
- wetb_temp: Wet-bulb temperature is the lowest temperature that can be obtained by evaporating water into the air. It measures the humidity of the air. Unit=1 deg C to the nearest 0.1 deg C
- stn_pres: Station air pressure. Atmospheric pressure as measured at the station level. Correction for altitude is not applied. Unit=1 hpa to the nearest 0.1 hpa.

- alt_pres: Altimeter pressure. Unit=1 hpa to the nearest 0.1 hpa.
- wmo_hr_sun_dur: This gives the readings from the newer automatic sun sensor which has now replaced the Campbell Stokes Recorder.

Of course, the dataset consisted of 94 more weather parameters, which are not essential for this research project. Therefore, the data which is important from the perspective of the research project had to be scrapped from the original dataset. and other parameters had to neglect. Finally, the following set of parameters were selected from the MIDAS UK hourly weather observation data for the years 2011 to 2014,

- ob_time
- wind_speed
- dewpoint
- stn_pres
- msl_pressure
- wetb_temp
- air_temperature
- wind_direction
- visibility
- wmo_hr_sun_dur

The above parameters are considered for the feature analysis in chapter 7.

5.4. Data cleaning and manipulation

Following data manipulation is performed to make the final dataset which will be used for the machine learning and active learning analysis in this research project.

5.4.1. Discarding the ToU tariff experiment data from LCL dataset

The LCL project had a limited number of experiments during the year 2011. Only 21 HIGH tariff experiments were carried out during the period of the LCL trials. As the number of trials was not sufficient for training a forecasting model, the data regarding ToU tariff signals and corresponding ToU response data is discarded. Also, the research project aims at selection methods for 'informative' ToU tariff days, the use of ToU data from the LCL project would have biased the results of the selection criteria. Therefore, we will only use fixed tariff data from the LCL dataset, discarding all the other data.

5.4.2. Matching the date-time-stamp of two datasets

The data recorded in the LCL trials is half-hourly, whereas, the weather data is an hourly data. Therefore, we need to match the timestamps of two datasets to make use of two datasets under one data frame. We could either interpolate the weather data to generate the half-hourly observations or we can skip every other energy consumption observation recorded in LCL trials to match the hourly resolution of weather data. The latter option involves loss of data and potential information. Therefore, the former method is chosen and all the continuous weather parameters are interpolated to double the time resolution of the weather dataset. Now both the datasets (LCL and weather) have datetime indices with the same timestamps.

Column	Datatype	Description
Index	datetime index	Includes the date and time information with half hourly resolution
Fixed tariff consumption (LCL)	continuous	Contains fixed tariff data of LCL users consumption by the users (total 4,334 columns)
temperature	continuous	Interpolated AIR_TEMPERATURE parameter from MIDAS dataset
dew point	continous	Interpolated DEWPOINT parameter from MIDAS dataset
wind speed	continous	Interpolated WIND_SPEED parameter from MIDAS dataset
Station pressure	continuous	Interpolated STN_PRES parameter from MIDAS dataset
MSL pressure	continuous	Interpolated MSL_PRES parameter from MIDAS dataset
Wet bulb temperature	continuous	Interpolated WETB_TEMP parameter from MIDAS dataset
wind direction	Categorical	Interpolated WIND_DIRECTION parameter from MIDAS dataset
visibility	Continuous	Interpolated VISIBILITY parameter from MIDAS dataset
Sun duration	Continuous	Interpolated WMO_HR_SUN_DUR parameter from MIDAS dataset
Hour of day	Categorical	Value of hour of a day calculated from 00:00 (time)
Day of week	Categorical	Day of week calculated from Sunday to Saturday
Month	Categorical	Month of year
Season	Categorical	Summer, Autumn, Winter or Spring

Table 5.1: The final data table with the description about the data

5.4.3. Generating calendar features

The electricity consumption pattern of residential users is often periodic. Therefore, it is necessary to inspect the effect of calendar effect on the performance of a forecasting model. The daily aggregate electricity consumption pattern can be captured by hour-of-day categorical feature, which will show the hour of a day based on the datetime index of a data point. Similarly, a weekly pattern can be captured by the day-of-week categorical feature. The other categorical calendar features which are considered are month and season of the year.

5.5. Final dataset

Table 5.1 shows the final data that will be used in this research project. As mentioned earlier, weather data is interpolated to match the datetime index frequency of the electricity consumption data. This dataset has a resolution of half-hour. The index matching is performed to match the appropriate features with the fixed tariff data of LCL trial users.

5.6. Summary

This chapter has introduced the data that is used for the research project. The LCL data and MIDAS weather data, both required tedious data cleaning and manipulation. The MIDAS weather dataset had to be interpolated to match the index of LCL trials dataset. The table 5.1 shows the final dataset which includes individual electricity consumption observations of 4,334 fixed tariff users of LCL project, selected weather features from MIDAS dataset and generated calendar features. As mentioned in chapter 4 the research project considers the aggregate electricity consumption data of users for fixed as well as ToU tariff policies. As of now, ToU tariff data and corresponding electricity consumption data used in this thesis are not introduced. Next chapter (chapter 6) will introduce the data simulator, which is designed for simulating aggregate fixed and ToU tariff response of the residential users.

6
Demand Response Simulator

This chapter discusses the design and implementation of Demand Response Simulator. Firstly, the limitations of the LCL (Low Carbon London) dataset are discussed. Secondly, the choices of the simulator design are described. Then the simulator setup is described, explaining each block of the simulator. Finally, the output of the simulator is explained.

6.1. Introduction

The electricity consumption data for this project is mainly based on the LCL project. The project gathered real-time electricity consumption data for over 5000 trial participants. This thesis tries to analyse the problem of forecasting electricity consumption under ToU tariff policy from energy retailer's point of view. Therefore, we are only concerned with the aggregate behaviour of users rather than their individual consumption patterns. When working with aggregate data, we have to consider the advantages and disadvantages of aggregating the data.

One of the major advantages of considering aggregate data is that combining the group of measurements of energy users reduces the individual variance of energy consumption patterns. If we consider an individual household for energy consumption prediction, then the variability in the user consumption behaviour is too complex to be fully understood by a machine learning model. For example, if a residential user may start charging their electric vehicle at any time between 19:00 to 22:00 (time), this randomness of switching on such a large load can not be learned by a demand forecasting model. But when we combine energy consumption of many such households, the variability in the consumption reduces. This reduces the input noise or randomness in the data. The clear trend in the data can be observed on the aggregate level. Therefore, forecaster learns concept with higher accuracy and with fewer training samples.

The forecasting of individual energy consumption pattern does not get affected by a total number of users. But when the number of participating users increase in the experiment, the variance of aggregate consumption data reduces, improving the trend in the data. Therefore, for large scale measurements of user energy consumption, forecasting of aggregate energy consumption is more accurate than the summation of forecasting for individual user energy consumption [14].

One downside of using the aggregate data is that it is hard to assess the energy consumption pattern of individual users. The individual assessment of the user consumption may help utilities to understand trends in the user consumption pattern. Also, using appliance detection methods [16], it is possible to calculate the per appliance consumption and provide energy consumption reports to users to help them improve their demand response, in this case, for the ToU tariff signals. This non-intrusive type of demand response scheme is out of the scope of this research project.

The following section discusses the problems associated with a direct user of LCL data for the experiment setup of this thesis.

6.2. Limitations of LCL dataset

The energy consumption data collected in the LCL project was collected for four years (2011-14). Out of those four years, the ToU tariff experiments were carried out only in one year (the year 2013). For the remaining three years, the irregularity in the data collection can be observed from the final dataset provided by the LCL authority. Especially, in the years 2011 and 2012, we observe sparsity in the user energy consumption data, which is not favourable for understanding the global user behaviour in the LCL trials setting. The LCL trials took place with 5,533 London electricity users. ToU experiments were carried out with only 20% of the total participants. All these conditions impose limitations on the quality of information that the dataset can provide for this project. The main implications of the small trial period (one year) are given below:

1. Seasonality, Trends of time series
2. The novelty effect
3. Limited ToU trials in LCL project

6.2.1. Seasonality, Trends of time series

Seasonality and trends are important characteristics of time series data. Seasonality can be defined as the linear or non-linear component that changes over time and repeats periodically. For energy usage time series data, the most visible seasonal factors are weekly seasonality and annual seasonal change (winter, spring etc.). The figure 6.1 shows the data of aggregate energy consumption of normal tariff users in LCL project collected from 2012 to 2014.

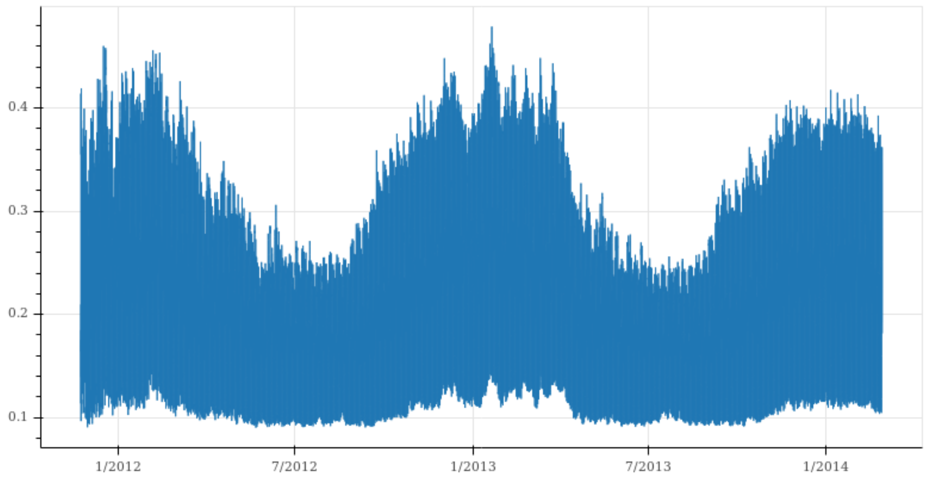

Figure 6.1: Seasonality of energy consumption pattern for residential users of LCL trials

The annual periodic changing pattern in energy consumption can be observed clearly. Usually, during summers the overall energy consumption drops due to the rise in ambient temperature. Whereas, during winters overall energy consumption increases due to falling in ambient temperature. These are the findings limited to the UK based on data and overall weather conditions. This user behaviour can be different for different regions/ countries based on local conditions. For this particular dataset, the ToU trials were conducted for one year. It will be a risky assumption to consider this data to be information-rich about user behaviour characteristics. In other words, even if the data captured the annual seasonality of energy usage pattern, there is no way to cross-validate that claim.

Similarly, the trial period is insufficient to understand any trends in the data. In time series, the trends are linear or non-linear components which do not repeat. In energy consumption data, trends can be listed as an overall increase in energy consumption over the years due to digitalization, the rise of electric vehicles and increased electric heating systems. Another example of a non-linear trend is irregular load demand due to increasing rooftop solar installations.

6.2.2. novelty effect

The novelty effect is mostly seen in the introduction of a new system or a new technology. The novelty effect represents a set of human traits which are motivated to act based on little or no experience with the system which often results in some bias in the initial phases of the testing. In the support of this argument, the original findings of LCL learning report say, ”the performance is consistent with a slight reduction in demand response magnitude for the trial”[24]. The participants of the LCL trials were willingly participating in the project. The desire for active participation might have resulted in biased outcomes during the initial weeks of trials. Therefore, using a machine learning model to learn user behaviour based on the dynamic tariff behaviour of users might not reflect actual response when implemented on a wide scale.

6.2.3. Limited ToU trials in LCL project

The Low Carbon London project carried out a limited number of ToU tariff experiments in one year (total of 21 HIGH tariff ToU experiment days). The given number of ToU tariff experiments are not enough to train a forecasting model with acceptable accuracy. Also, the purpose of active learning in this thesis is to investigate the methods for selective sampling of days for ToU tariff experiments, which can not be fulfilled by the use of existing data of HIGH tariff ToU experiments. The use of this data may bias the active learning algorithm, therefore the use of this data is avoided.

All the evidence presented here suggests that a lack of adequate data would cause a problem in implementing the machine learning framework. The two-fold problem has to be addressed before implementing the machine learning algorithm - firstly the issue related to limited data has to be resolved, and secondly, the inconsistency in the demand response by the users (novelty effect) has to be rectified.

6.3. Generating new ToU tariff trials

We will discuss the ways of creating additional ToU trials strictly from the thesis perspective i.e. with very limited resources, time and budget etc. There are two straightforward ways of achieving the goal of increasing the number of ToU tariff trials- 1) Perform trials in a real environment, and 2) Simulating trials in the simulated environment. In an ideal scenario, trials in a real environment would produce the most accurate and reliable results for ToU tariff signals. But this is not the most cost-effective solution. As discussed in previous chapters, setting up an experiment of this size requires tremendous resources, time and budget.

This thesis approaches the problem of limited data by creating data simulator. Conventional simulators are based on little to no data, therefore makes the validation of the simulator difficult. Instead, the proposed simulator is a data-driven simulator. The simulator utilises the fixed tariff electricity consumption [1] data gathered during the LCL trials. The following section discusses the simulator setup designed for the experiment.

6.4. Simulator setup

The simulator aims to generate ‘new’ ToU and fixed tariff days which would imitate real-world conditions as close as possible. As the ‘new’ data is a processed subset of actual observations, the effect of weather conditions, seasonality and various calendar effects such as weekends, holidays etc are already covered by the original observations. Manually modelling with these parameters would require highly complex simulator design. Therefore, in this project manual modelling of electricity consumption is avoided.

6.4.1. Generation of synthetic fixed tariff day instance

One of the common ways of generating a new instance of data is by adding noise in the original data. Adding white noise will be one of the simplest implementations of this method. But, white noise may erase some important information from the energy consumption data and will lead to a synthetic consumption data which is not representative of real-world consumption. Therefore, synthetic data generation by addition of noise is avoided.

Instead, we will take a subset of fixed tariff electricity users and treat it as a different data point. We will randomly select a different set of users for calculating the average aggregate electricity consumption data. As we are only concerned with the average value of aggregate consumption, instead of using the full set of user, we will only consider 70% users for every new data point. By randomly selecting different users for each round of data generation, we will have a different combination of electricity consumption pattern which will lead to slight variation in the aggregate consumption. Therefore, we can have multiple fixed tariff electricity consumption patterns for the same day. Of course, each of the

[1]The LCL project recorded electricity consumption of around 4,334 users who were subjected to fixed tariff

newly generated aggregate consumption half-hourly energy consumption sequence will correspond to the same weather and calendar conditions. Therefore, in a nutshell, we generate multiple instances of a day by aggregating a random subset of individual user energy consumption.

Algorithm 1 User selection for calculating aggregate fixed tariff consumption

```
procedure Random User Selection(C_n)
    Randomly select a day from the pool of LCL data C
    Collect the consumption data of n users for that day c_n
    Randomly select u users, where u = n * 0.7
    Calculate y_fixed = (1/u) Σ_{i=1}^{u} c_i
    Return y_fixed
```

The process of generating new aggregate user consumption instance is shown in algorithm 1. The input C_n is the pool of actual fixed tariff electricity consumption data of n users from LCL dataset. The algorithm outputs y_{fixed} which is an average of aggregate electricity consumption pattern of a day under flat tariff. The c_i is the 24 hours individual electricity consumption of ith user under flat tariff. As the energy consumption data was observed half-hourly under LCL trials, y_{fixed} and c_i both are 48 element vector.

6.4.2. Generation of synthetic ToU tariff day instance

The algorithm 1 only shows the process of generation of fixed tariff electricity consumption for a day. The process is comparatively straightforward as we are just aggregating the actual electricity consumption of randomly selected users for a day from the given input dataset. As the dataset does not consist of ToU tariff days, now we need to create a mechanism to synthesize a demand response for a ToU tariff policy. Before proceeding to the actual calculation of demand response, we need to design a tariff policy.

The tariff policy

The overall design considerations of tariff policy are shown in chapter 4. To recap, only two tariff levels are considered for the research project - 1) NORMAL tariff, and 2) HIGH tariff. We will focus on reducing the peaks of residential energy consumption. Therefore, the HIGH tariffs will be sent only for the duration of the peak energy consumption. Otherwise, NORMAL tariff will be sent for all the other half-hourly time slots. For each new policy starting and ending of the peak hours are randomly chosen from the set of available timings. Following equation shows the tariff policy

$$T(\tau) = \begin{cases} 1, & \text{if } \mathcal{X}_{\text{lower}} > \tau > \mathcal{X}_{\text{higher}} \\ 0, & \text{if } 0 > \tau > \mathcal{X}_{\text{lower}} \text{ OR } \mathcal{X}_{\text{higher}} > \tau > 48 \end{cases} \tag{6.1}$$

Where, $\mathcal{X}_{\text{lower}} \sim U(33, 37)$ represent the starting of the HIGH tariff signal (randomly selecting time between 16:30 (33rd slot of a day) and 18:30 (37th slot of a day) from discrete uniform distribution) and $\mathcal{X}_{\text{higher}} \sim U(43, 47)$ represent the ending of the HIGH tariff signal (randomly selecting time between 21:30 (43rd slot of a day) and 23:00 (47th slot of a day) from discrete uniform distribution) for a 48-element vector representing half-hourly tariff policy. As the tariff policy design is now explained, let us consider the modelling of user behaviour for any given tariff policy.

Modelling of user behaviour

The research project considers a single step change in the individual electricity consumption as a response to the change in tariff level. For example, the participating users of the ToU tariff experiment will reduce their loads by a single step (a predefined fraction of their potential fixed tariff consumption) when HIGH tariff time block of a day starts. But if all the users start their demand response at the same time, the aggregate demand response will be fairly mechanical (and may not reflect the real-world behaviour).

To model user behaviour, the users' load switching latency is considered. This can be explained by the following example. Let us say, for a particular day, the high tariff prices will start from time 17:00. In the real world, very few users will start reacting to the prices exactly at 17:00 o'clock. In most of the cases, people will start the load switching operations either slightly before or after 17:00 o'clock.

This behaviour is modelled by dynamically assigning the start and stop time of demand response to individual users for the given tariff policy. As the simulator generates demand response only for HIGH tariff time slots, the unique tariff signals for each individual user is calculated by,

$$T_{\text{user}}(\tau)_i = \begin{cases} 1, & \text{if } (\mathcal{X}_{\text{lower}} + \delta_i) > \tau > (\mathcal{X}_{\text{higher}} + \delta_i) \\ 0, & \text{if } 0 > \tau > (\mathcal{X}_{\text{lower}} + \delta_i) \text{ OR } (\mathcal{X}_{\text{higher}} + \delta_i) > \tau > 48 \end{cases} \tag{6.2}$$

Where, $\delta_i \sim U(-3,3)$ is the user latency variable, where δ_i is drawn from a discrete uniform distribution to generate tariff policy profile $T(\tau)_i$ for user i. The purpose of the dynamic allocation of demand response is to mimic the real-world behaviour of users and also to increase the complexity of the demand response behaviour.

Calculation of ToU tariff response

In the design of this simulator, it is considered that only a fraction of residential users will participate in the ToU tariff response while other users will continue with their usual electricity consumption pattern as if it is a fixed tariff day. The consideration was inspired by the availability factor of the users. For a given ToU experiment, some users may not be able to reduce their loads as they may not be available for the HIGH tariff time period. Also, some users may not want to suffer through inconvenience on a particular day. So, it is assumed that 70% users will participate in the ToU experiment response, while other 30% will continue with their usual electricity consumption.

Finally, the demand under ToU tariff policy is calculated by following the findings of LCL trials. LCL trials found that the average peak load shaving of 8% is observed during demand response for HIGH tariff signals[24]. We assume that out of n total users, only u number of users will participate in the ToU tariff response, where $u = n * 0.7$. Therefore, when c_i is the actual fixed tariff consumption of ith user from LCL dataset, we calculate the 24 hour ToU tariff electricity consumption y_{tou} as,

$$y_{\text{tou}} = \frac{1}{2}\left(\frac{1}{u}\sum_{i=1}^{u} c_i * T_{\text{user}}(\tau)_i * (0.88) + \frac{1}{n-u}\sum_{i=u+1}^{n} c_i\right) \tag{6.3}$$

Where, $T(\tau)_i$ is a 1×48 tariff policy profile for user i. When 70% of users contribute to 8% of total demand response, other 30% users will continue their normal fixed tariff electricity usage. Therefore, to achieve an 8% overall demand response, 70% of the users will effectively have to contribute around 12% in the demand response. This way, the above equation is used to calculate the average aggregate electricity consumption under ToU tariff policy. Next, we will look at the functional block diagram of the complete process.

6.5. Final data generation process

Figure 6.2 shows the block diagram of the data generator design. The following section provides information about the final dataset generation.

- As discussed previously, simulator accepts inputs of LCL fixed tariff electricity consumption data for years 2013-14 and corresponding weather and calendar features F.
- A random date is selected from the datetime indices of the input data and the consumption data C for the selected date is chosen for further process.
- Random subset of users is selected for generation of synthetic aggregate fixed tariff and ToU tariff consumption.
- The fixed tariff consumption y_{fixed} is calculated by aggregating the actual fixed tariff response of the above subset of users.
- A random base tariff policy $T(\tau)_i$ is calculated with equation 6.1 The tariff policy profile for participating users is then calculated as shown in equation 6.2.

- Based on the policy profile of participating user, ToU tariff response of each participating user is calculated. The Aggregate ToU tariff electricity consumption y_{tou} is calculated as shown in equation 6.3.
- The fixed tariff electricity consumption y_{fixed}, ToU tariff policy $T(\tau)_i$ and corresponding ToU tariff electricity consumption y_{tou} is then combined with features f which correspond to the selected date. This 24-hour dataset (48 half-hourly values for each variable mentioned above) is the output of the simulator.

This simulation process takes a long time to simulate the data for one synthetic day. Therefore, to reduce the time for execution of each active learning model, a pool of synthetic data is generated beforehand. The above process is repeated to simulate 10,000 synthetic days. The output of the simulator after the generation of each new day is stored in a dataset. The final dataset has the output data of 10,000 simulated days.

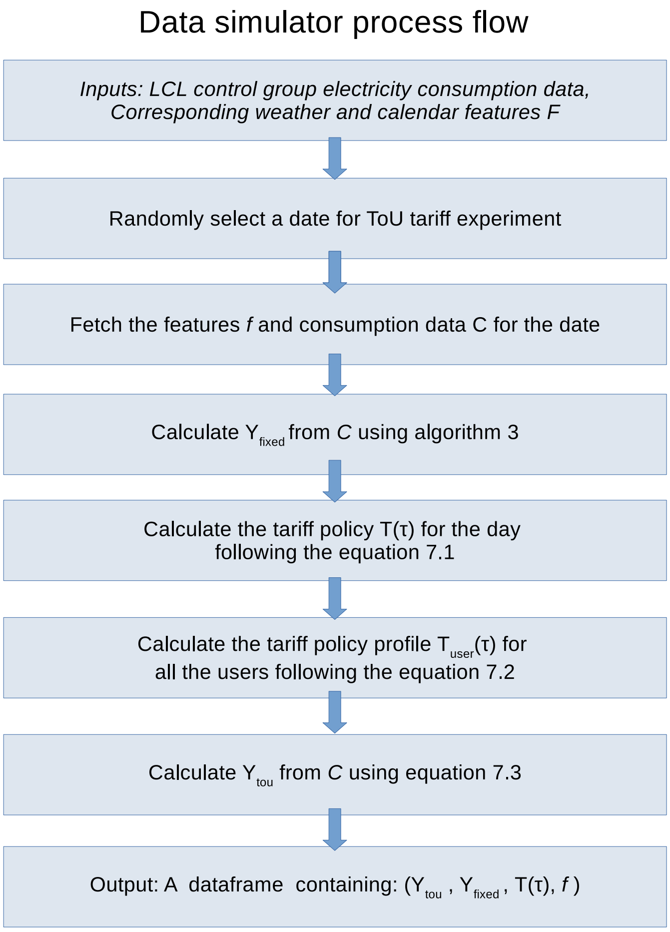

Figure 6.2: Block diagram of data simulator design

6.6. Discussion on Simulator output

- The output of the simulator is a 24-hour timeseries containing the energy consumption data along with other input features for a selected day. The simulator synthesizes the following variables as part of the output,
  - Fixed tariff electricity consumption y_{fixed}
  - ToU tariff policy $T(\tau)$
  - ToU tariff electricity consumption y_{tou}

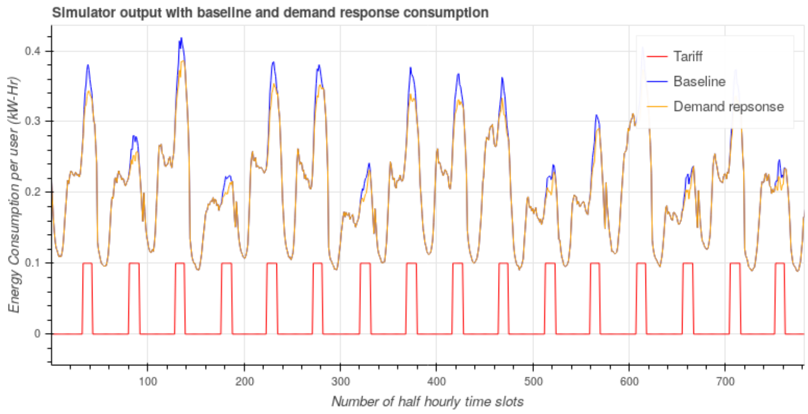

Figure 6.3: Baseline consumption and corresponding demand response consumption generated by the simulator

- The features for the model are classified as categorical and continuous variables. Continuous variables are numeric variables that have an infinite number of values between any two values. Most of the physical measurements of entities are continuous e.g. Temperature, Energy consumption etc. Whereas, categorical variables contain a finite number of categories or distinct groups. Categorical data might not have a logical order. The detailed information about the features and criteria for feature selection is shown in chapter 7. The selected features are shown in table 6.1.

Feature name	Feature type
Temperature	continuous
Dew point	continuous
Pressure	continuous
Wind speed	continuous
Sun duration	continuous
Day-of-week	categorical
Season	categorical
Hour-of-day	categorical
Month	categorical

Table 6.1: The features for the synthetic electricity consumption data

  These features are then correctly indexed with the data. Mind that, as we are synthesizing multiple instances of a day from the limited actual data, a day may be repeated. In that case, multiple instances will have the same input features as the underlying actual electricity consumption data correspond to the same input conditions. Table 6.2 shows the contents of the final data frame object output with their properties.

- The graphical timeseries illustration of baseline and demand response consumption is shown in figure 6.3. The consumption pattern for 16 consecutive days is plotted in the figure. The X-axis corresponds to the half-hourly time slots (no unit), and Y-axis corresponds to the average half-hourly energy consumption per user in (kW-Hr per half hour). The fixed and ToU tariff response energy consumption is plotted against the number of half-hourly time slots. Whereas, tariff signal, being an unit-less binary variable, is plotted against the half-hourly time slots to show the relationship between the demand response and the dynamic tariff signals.

  From the plot, we can see that the demand response consumption deviates from baseline consumption only when the HIGH tariff signal is provided. This is the desired output for a demand response experiment. The tariff response smoothing effect can also be seen at either end of the HIGH tariff period. The highest demand response (highest peak shaving) is observed mostly in the middle of the HIGH tariff signal. Secondly, in some cases, the midpoint of HIGH tariff signal and peak of the baseline consumption do not match. This will create an interesting case for the training of machine learning model, reducing the chances of overfitting the data.

Variable name	Data type
Fixed tariff consumption	continuous
ToU tariff consumption	continuous
Tariff policy	categorical
Temperature	continuous
Dew point	continuous
Pressure	continuous
Wind speed	continuous
Sun duration	continuous
Day-of-week	categorical
Season	categorical
Hour-of-day	categorical
Month	categorical

Table 6.2: Contents of simulator output dataframe object. Each output simulator output contains the 48 entries of the variables shown in above table for the a base day measured with the half hourly resolution.

6.7. Summary

The simulator provides an output which is based on the real-world energy consumption data. Therefore, unlike a theoretical model which is based on a set of equations, the generated energy consumption profile is closer to the real-world energy consumption. The simulator is used to generate the data for testing the algorithms presented in chapters 7, 8.

7

Basic forecasting and analysis

This chapter introduces the implementation of the demand response forecasting model. This chapter provides an introduction to the experimental setup of the considered demand response problem. The feature selection process for the model is then explained. This chapter explains the forecasting model used for predicting consumption of residential users under demand response of ToU tariff. The results of the implementation are discussed at the end of the chapter.

7.1. Introduction

The demand profile forecasting for under the fixed tariff and ToU tariff is considered for generating a baseline model. This forecasting model requires the knowledge of historic target parameters. In this case, the target variable is the half-hourly measurement of the average consumption of electricity (measured in kW-hr per half hour) by the particular set of users as a response to dynamic tariff policy. The target variable depends on the feature variables which partially contain information about the target variable. These features do not usually have one to one correlation between the target variable and other individual feature variables. This chapter will first focus on feature selection and analysis, followed by the machine learning implementation.

7.2. Feature selection

Timeseries regression data contains multiple observations of target variables and corresponding features. It is important to choose the right features which can map the input space on the output target variable of the training data. Following features are considered for implementing the forecasting model of the demand profile.

- Weather dependent features: Weather parameters like temperature, dew point, pressure, wind speed etc.
- Calendar dependent features: Day of week (S-M-T-W-T-F-S), Season of the year (Summer, Autumn, Winter & Spring), Hour of the day, month of year
- Tariff for the residential electricity

The reasoning behind the selection of the above features is given below:

- The weather parameters highly affect the pattern of residential energy consumption [21]. Usually, this effect is aligned with the comfort factor of residential energy usage. Usually, users tend to change their energy usage pattern to adapt to surrounding ambient weather.

    1. Most of the heating and cooling load consumption patterns have a high correlation with ambient temperature. People tend to increase heating consumption during low temperatures (usually, in winter) to maintain a comfortable temperature inside the closed walls of a building. Similarly, during summers, the energy consumption is affected by higher outside

temperature as people seek comfort in the cooler air of air conditioning. As heating and cooling loads are one of the largest loads of residential users, the temperature is an important indicator of the level of residential energy consumption. Refer to figure 7.1 for visualizing the correlation between temperature and user energy consumption. The timeseries in the plot are the scaled differences between the actual value and the rolling mean of the data. Subtracting the rolling mean from the actual value improves the seasonality and removes the trends from the timeseries. Scaling the data makes it easier to compare the two timeseries.

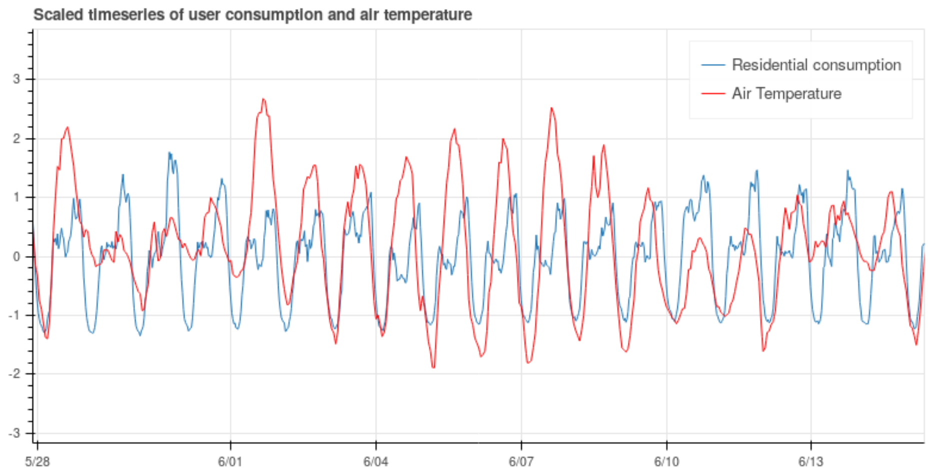

Figure 7.1: Correlation between temperature and and user electricity consumption. The plotted data is a difference between the actual value and the rolling mean of the data

2. The dew point is the temperature to which air must be cooled to become saturated with water vapour. When further cooled, the airborne water vapour will condense to form liquid water. When air cools to its dew point through contact with a surface that is colder than the air, water will condense on the surface. The dew point is a good indicator of humidity of the air. During the rainy weather often high humidity is observed. Humans perceive the temperature differently under different levels of humidity. Therefore, the dew point is an interesting factor to put under research while modelling the users' electricity consumption profile.

3. The dew point, pressure and wind speed create a 'chill factor' which changes the perceived temperature for the users.

- The energy consumption of residential users is highly dependent on calendar events. For example, from the data it is observed that there is a significant difference in the peaks of weekdays and weekends. Peak hours are usually observed during the evening times. Other than the 'day of week' calendar factor, hour of day affects the energy consumption pattern. For example, during the nights, energy consumption is often reduced as that time period coincides with 'sleep time' of the users. Other important calendar features are month and season of year, which incorporate the yearly seasonal effects.

- Unlike the other two categories of the feature variables, electricity tariff is a human controllable variable. The dynamic tariff effect on user energy consumption does not need any new introduction. The simplest form of dynamic ToU tariff policy includes two tariff levels - NORMAL and HIGH, for off-peak and peak time periods respectively. Customers try to reduce their consumption during peak time blocks to reduce electricity costs. In all the other time blocks, the users will continue their usual pattern of energy consumption. Therefore, it is interesting to analyse the relationship between the tariff signals and user consumption pattern during those tariff signals. This research project only considers the binary ToU tariff case as described in chapter 4.

7.2.1. Feature correlation

All the parameters satisfying above feature selection criteria are considered to calculate the correlation plot for the features. Correlation is used as a method for feature selection and is usually calculated between a feature and the output class (filter methods for feature selection). Pearson correlation is one of the oldest correlation coefficients developed to measure and quantify the similarity between two variables [5]. Formally, Pearson's correlation coefficient is the covariance of the two variables divided by the product of their standard deviations. Pearson correlation has the following limitations -

- Strong influence of outliers — Pearson is quite sensitive to outliers
- Assumption of linearity — The variables should be linearly related, therefore it does not work with the categorical variables
- Assumption of homoscedasticity, meaning each random variable in the sequence has the same finite variance.

The Pearson correlation makes the above assumptions. Therefore, it has limited applications where the data is noisy or non-linear. But Pearson correlation provided some early insights about the data, such as -

- Correlation can help to predict one variable from another variable (in case of missing values)
- Sometimes, correlation can indicate the presence of a causal relationship

Let us have a closer look at the various correlation types -

- Positive Correlation: It means that if feature A increases then feature B also increases or if feature A decreases then feature B also decreases. In positive correlation, the variables have a linear relationship.
- Negative Correlation: means that if feature A increases then feature B decreases and vice versa.
- No correlation

The correlation plot between the features used for current implementation is shown in figure 7.2. For the features, the absolute correlation value is showed in the plot. The lighter shade represents a high correlation between the features, whereas darker shades of blue indicate a lower correlation between the features. Where the names and details of the features are explained in chapter 5.

We can see a very high correlation between the average energy consumption data (namely, 'expected' & 'response'). This correlation is obvious as the two energy consumption data are essentially very similar to each other; As the considered energy consumption observations for any day is equal for NORMAL tariff signal, it was expected to observe high correlation between We can also see the correlation of energy consumption data with temperature ('AIR_TEMPERATURE'), dew point ('DEWPOINT'), air pressure ('MSL_PRESSURE') and wind speed ('WIND_SPEED'). We will analyse these features into details as it is observed that the electricity consumption of residential users is a function of weather patterns [19]. From the categorical features, we will keep hour of the day ('hod') feature as from figure 7.1 shows that the there is a high correlation between electricity consumption and time of a day. We will also keep tariff signals as a feature and we know that demand response under ToU tariff policy is a function of tariff signals.

Now we will look at the simple forecasting model for energy consumption at the fixed tariff.

7.3. Forecasting of fixed tariff energy consumption

The forecasting of timeseries regression is a supervised machine learning process. In this section, we will look at the basic machine learning model using random forest regression algorithm. The model will be trained on the features selected in section 7.2.1 to predict the average energy consumption of the users. We will only consider fixed tariff energy consumption for the problem. The following section discusses the setting for the regression model.

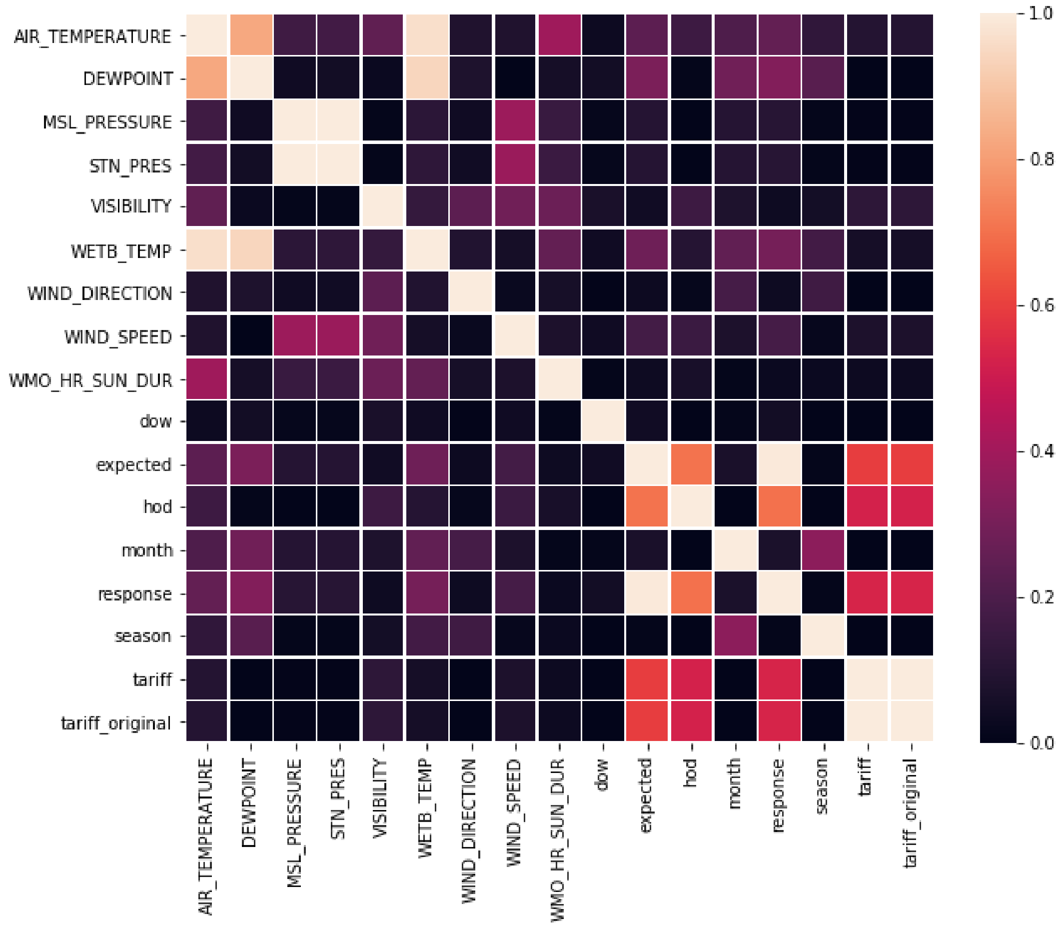

Figure 7.2: Data correlation plot (absolute correlation values are considered)

7.3.1. Problem setting

In this model, the aim is to implement the random forest regression model for predicting the average energy consumption of users for the fixed tariff conditions. The basic working principles of the random forest regression model is explained in chapter 3. As we will be predicting the future electricity demand for ToU tariff policy, the weather inputs are usually predictions of the future weather. Of course, the quality of the predictions of our forecasting model depends on the quality of the weather predictions. But considering the scope of this project, we will consider that the weather forecasts are accurate and base our results on the assumption. As an implementation detail, the actual electricity consumption data used in the project is historic data and therefore, the weather parameters used in this thesis are actual values of historic observations. The six inputs to the machine learning model are shown below:

- Weather inputs:
  1. Air temperature (°C)
  2. Dew point (°C)
  3. Pressure
  4. Wind speed (knots)
- Calendar inputs:
  1. Hour of day
  2. Day of week

As explained in chapter 5, like all the other data, these inputs have half-hourly measurements. Mind that, as the energy consumption data is 'measured' (simulated) under fixed tariff we will have a fixed tariff for the entire data. Therefore, we are neglecting the tariff from potential feature list for the current model.

The output of the model (target variable) is average energy consumption (kWh per half hour). The target variable is of size 1×1. Therefore, using random forest regression model we try to map $1 \times n$ dimensional data of n features to every target variable of size 1×1. As we are considering 6 features for the specified model, we have $n = 6$. Therefore, we will predict a 1×1 target variable for every input of $1 \times n$ input feature variables.

7.3.2. Evaluation criteria

The evaluation criteria are required to evaluate the performance of a machine learning model. The quality of the forecast can be assessed by the evaluation criteria. For this research project, we consider the Mean Squared Error (MSE) as the performance evaluation criteria for the following machine learning models. When $\hat{Y}_q$ the vector q predictions are generated from q data samples for the observed values indicated by Y_q, then the MSE for the above setting is defined as:

$$\text{MSE} = \frac{1}{q}\sum_{i=1}^{q}(Y_i - \hat{Y}_i)^2$$

MSE is simple to calculate and the analysis of the model over MSE is easy to understand. The calculation of MSE requires actual observations.

Secondly, in the dynamic tariff scenarios, the experiments of a ToU trial are compared using different budgeting cases. For example, we will analyze how many fixed tariff days can bring the same improvement in the energy consumption forecasting model which is brought by introducing a ToU tariff experiment day.

7.3.3. Random forest model

The random forest model is considered with the following parameters. The number of estimators is tuned and kept at 10. The mean squared error for the different number of training samples is calculated and convergence of the model is checked. The previously mentioned feature set along with observed target variable is used to train the random forest model. As a day has a fixed number of data samples (48 samples), now onward, the data is measured in terms of 'days'. Therefore, a day sample is defined as 48 electricity consumption observations (fixed & ToU tariff consumption) along with the corresponding features of that day. The test data set of 500 fixed tariff days is kept constant throughout the process.

Figure 7.3 shows the predicted energy consumption of users along with the observed energy consumption for the given input parameters. The plot is generated by the test dataset of 500 fixed tariff days tested on the model with the training dataset of 640 fixed tariff days.

The machine learning model with the same hyperparameters is trained on the following number of fixed tariff training days - [20, 40, 80, 160, 320, 640, 1280, 2560, 5120, 8000]. The training sample days are randomly selected from the data pool generated by the simulator shown in chapter 6. Similarly testing dataset is randomly selected before training the model and kept the same for one round of training session (from 20 days to 8000 days). The average of MSE of 5 such training sessions is calculated and plotted against the number of training days. The log-log plot of MSE over the various number of training samples is shown in figure 7.4.

From figure 7.4 it is clear that increasing the number of training samples reduces the mean squared error of predictions. This proves that the performance of the machine learning model increases with the increase in the number of training samples. Increasing the number of training samples reduces the MSE of the model further, indicating that the model can still learn with new training data. Obtaining the absolute convergence requires a large amount of data which is not possible under the scope of this thesis.

The post-training feature analysis is performed by plotting the feature importance of the machine learning model with 5120 training days. The corresponding plot is shown in figure 7.5. We see the similarity in the data correlation matrix shown in figure 7.2 and feature importance plot drawn from

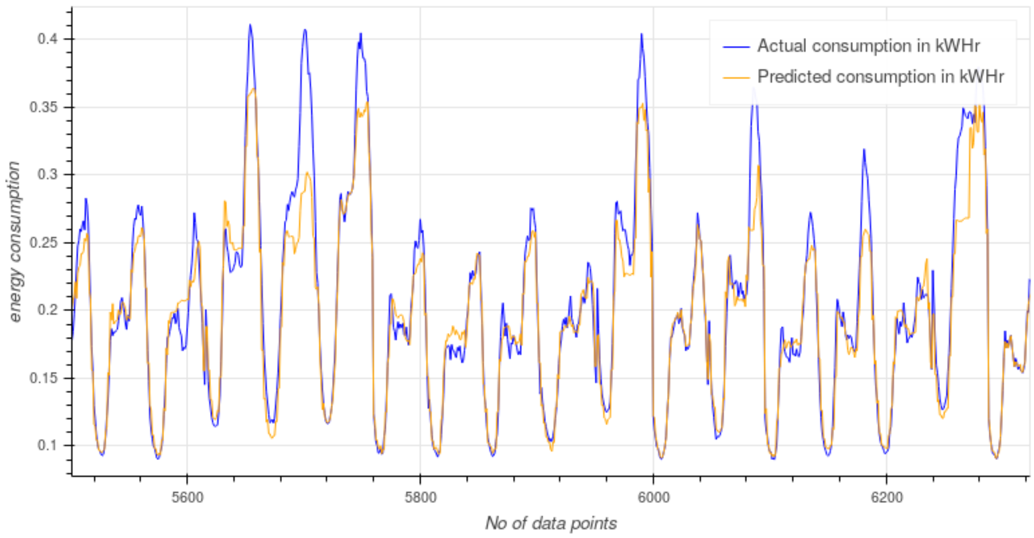

Figure 7.3: Predicted average fixed tariff energy consumption of users and observed actual consumption at 100 fixed tariff days training size

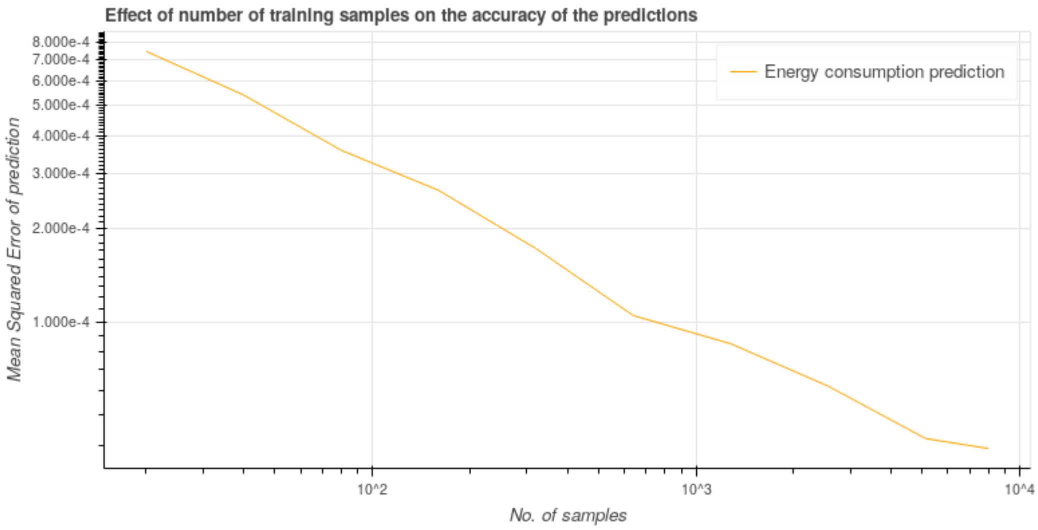

Figure 7.4: Log-log plot of MSE over the number of training samples

the trained model shown in figure 7.5. Therefore we intend to keep the previously selected features for further analysis of the problems and implementation of next algorithms.

Next, we will consider the dynamic tariff case. The following section covers the implementation and results of the forecasting of the average energy consumption of users under ToU tariff prices.

7.4. Forecasting of ToU tariff energy consumption

For forecasting ToU tariff response, we need to modify the previous problem setting to incorporate the dynamic tariff and the energy consumption related to the dynamic ToU tariff. As mentioned earlier, the ToU tariff is binary including NORMAL and HIGH signals. The revised inputs are given as:

- Weather inputs:
  1. Air temperature (°C)
  2. Dew point (°C)
  3. Pressure

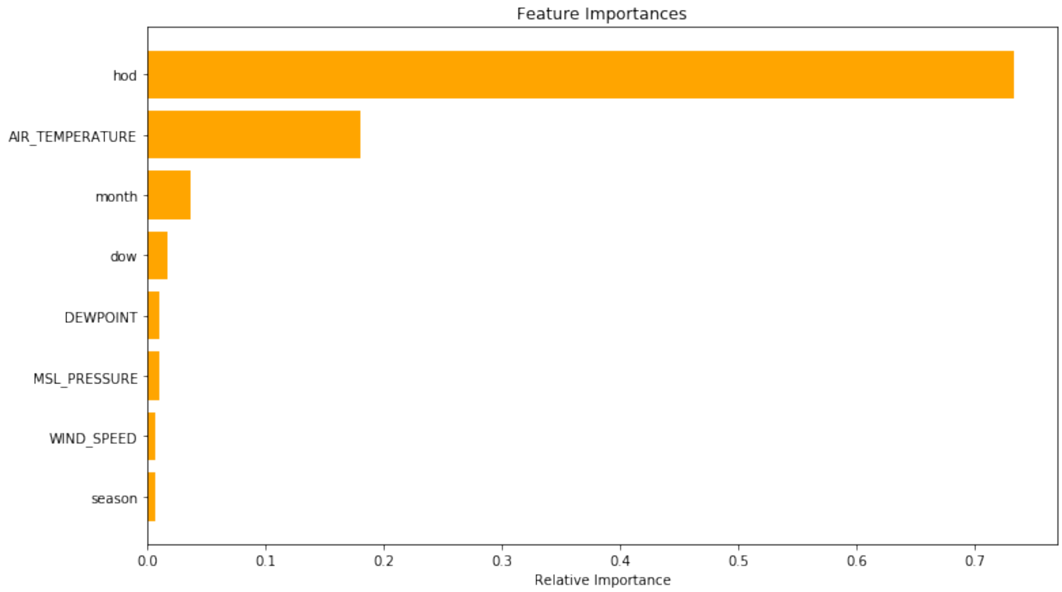

Figure 7.5: Feature importance plot for random forest regression model with 5120 training days

4. Wind speed (m/s)

- Calendar inputs:
    1. Hour of day
    2. Day of week
- half-hourly ToU Tariff policy for every day

Similar to the last case, each of the variables is available for half-hourly observations of average electricity consumption. Therefore, data corresponding to each day includes 48 values of each input variable. Now, with the introduction of ToU tariff as a feature, we want to predict the demand response under the ToU tariff policy. The output target variable is of size 1×1 and corresponding input feature variables are of size $1 \times n$, where n is several features which is equal to 7.

In this section, we will consider three models-

1. Model 1 trained on ONLY fixed tariff energy consumption observations
2. Model 2 trained on ONLY dynamic ToU tariff energy consumption observations
3. Model 3 trained on fixed AND dynamic ToU tariff consumption observations (30% dynamic tariff observations, 70% fixed tariff observations)

For the first two cases, the models will be purely trained either on the fixed tariff energy consumption or on the dynamic ToU tariff energy consumption. The third model, however, is partially trained on both types of energy consumption data. As shown in the input feature list (presented above), the tariff is used as input for all three models. The tariff corresponding to the fixed tariff consumption will only be considered at 'NORMAL' level. For dynamic ToU tariffs, the 'HIGH' level tariffs will be provided during the peak time blocks an experiment day. These 24-hour tariff policies are pre-determined by the simulated energy retailer (chapter 6). The more information on the design of tariff policies is provided in chapter 4.

These three models are then tested on a dataset of dynamic ToU tariff experiment days. That is the performance of these models evaluated for predicting the energy consumption during demand response for the given ToU tariff policy. Mind that, for the prediction of energy consumption under ToU tariff policy, we provide all the input feature variables including the ToU tariff policy for the subjected time blocks. Figure 7.6 shows the comparison plot between the three models. The machine learning model with the same hyperparameters is trained on the following number of training days - [20, 40, 80, 160, 320, 640, 1280, 2560, 5120, 8000]. The training sample days are randomly selected from the data pool generated by the simulator shown in chapter 6. Similarly testing dataset is randomly selected before training the model and kept the same for one round of training session (from 20 days to 8000 days). The average of MSE of 5 such training sessions is calculated and plotted against the number of training days. Again, the MSE of the predictions and the observed target value is plotted against the number of training samples.

We can see that all three models tend to reduce the prediction error by increasing the total number of training samples; This effect is observed as by increasing the number of training samples, models also learn the basic energy consumption pattern of the average demand of users. As discussed in chapter 4, the customer behaviour does not alter for any other time blocks, except for the HIGH tariff time block. Therefore, even if a model is trained on the fixed tariff samples, the model predicts electricity consumption during off-peak time blocks with higher accuracy. Also, off-peak time block usually occupies up to 75%-80% of the day time in some cases of dynamic ToU tariff policies.

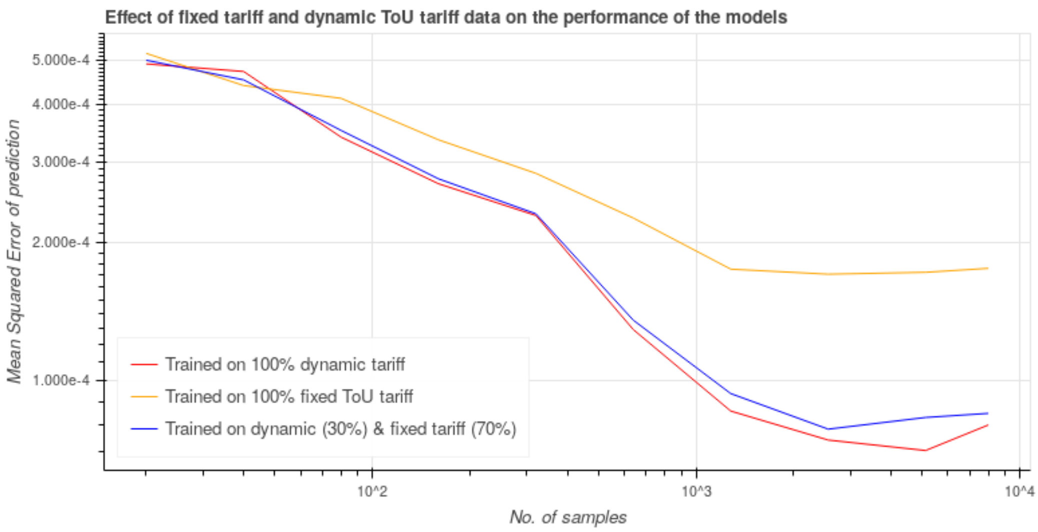

Figure 7.6: The comparison plot of the effect of fixed tariff vs dynamic ToU tariff energy consumption training data on the performance of the machine learning model.

We also see that the Model 1 (100% fixed tariff training samples) MSE performance settles on the noise bed much faster than the model shown in figure 7.4. We see that the model slightly improves with the increase in training samples. The model fails to achieve the prediction accuracy of the other two models. This indicates that the newly introduces demand response behaviour of the users (from test data) is outside the knowledge domain of the Model 1. But similar to the figure 7.4, the model can improve the forecast of a day by improving the accuracy over NORMAL tariff time blocks. But in Model 2 and Model 3, we see drastic improvement by the introduction of ToU tariff sample days. Another important observation can be made by comparing the Model 2 and Model 3. The difference between the MSE performance of Model 2 and Model 3 is quite low. Therefore, it can bee inferred that the introduction of dynamic ToU tariff experiment days in the training data samples improves the performance of a demand forecasting model. We see the improvement in the performance is assisted by the fixed tariff sample days.

If we assume that fixed tariff samples have no information (or negligible information) about the ToU tariff days, then Model 3 would require around 3 times more data points to achieve the same level of performance of Model 2 (Model 3 has 30% ToU tariff days); which is not the case. The performance of model 3 is quite similar to the Model 2.

From the above observations we can conclude the following things:

- The Model 1 improves the performance slightly before reaching the saturation, indicating that the lack of information about the demand response from fixed tariff samples results in a knowledge gap in the model.

- But fixed tariff samples help the Model 2 to improve the performance and achieve similar results as shown by Model 3. We can conclude that even if we have a small number of ToU tariff samples, we can improve the overall performance of the ToU tariff demand forecasting model by providing fixed tariff data.
- This does not necessarily mean that the Model 2 will perform equally well while predicting the peak electricity consumption under HIGH tariff signal.

7.4.1. Budgeting factor

Let us now consider the budgeting factor. Let us assume that energy retailers have low operation cost for the fixed tariff days and higher cost for the dynamic ToU tariff experiment days. We expect some improvement after implementation of the demand response trials. This setup requires a lot of resources, planning and operation costs. The assumption for the implementation of such projects is that this initial innovation cost will be covered by the profits gained by future ToU tariff schemes using the same infrastructure. The electricity retailer can buy electricity either in the spot market or in the intraday market. So if electricity retailer has poor performing demand forecasting model, they might need to buy electricity in the intraday market. As seen in the last section, the bottleneck of a forecasting model performance is the peak load forecasting under ToU tariff policy, the electricity retailer will need to buy electricity for peak hours in intraday, which is expensive. Therefore, having a better model will improve the profits of the electricity retailer. Direct financial profits are easier to calculate. But successful implementation of ToU tariff scheme can also reduce the cost of production and infrastructure in the long term. Therefore, the implementation of ToU tariff policy could be a financially profitable action for multiple electricity sector players.

The question then arises is that how much cost of an experiment is 'acceptable' considering the improvement in the demand forecasting model performance? Here, we would like to observe the number of fixed tariff data points required to attain the same level of performance for every dynamic ToU tariff experiment data point.

For this analysis we consider three models:

- Model A is introduced to only 10 ToU tariff days (one day sample contains 48 data points)
- Model B is introduced to only 50 ToU tariff days
- Model C is introduced to 100 ToU tariff days

If the above models are trained on 1000 day samples, it means that the training data contains (1000 (minus) dynamic ToU tariff days) number of fixed tariff days. But if we consider 20 training samples for all three models, the Model A will select 10 ToU tariff days and 10 fixed tariff days; whereas, other two models (Model B (50 ToU days) and Model C (100 ToU days)) will only consider ToU days to fulfil the requirement of 20 days samples. The machine learning model with the same hyperparameters is trained on the following number of training days - [20, 40, 80, 160, 320, 640, 1280, 2560, 5120, 8000]. The training sample days are randomly selected from the data pool generated by the simulator shown in chapter 6. Similarly testing dataset is randomly selected before training the model and kept the same for one round of training session (from 20 days to 8000 days). The average of MSE of 5 such training sessions is calculated and plotted against the number of training days. Figure 7.7 shows the performance of all three modes. The performance of Model B and Model C for smaller training data size is nearly equal.

Another observation is that Model A starts lagging in performance even for the smaller training samples. We can see the effect of dynamic ToU tariff days in this case. As for training set of 20 days, Models B and Model C contain all the observations from ToU tariff days; This provides the two models with an early improvement in the performance. We also see that Model A converges early to attain the equilibrium at the noise bed.

If we consider the performance of Model A at the saturation bed, it requires about 4200 more fixed tariff day observation to achieve the same performance (MSE = 10^{-4}) as of Model B and about 4400 more to achieve the performance of the Model C. For Model B to achieve the performance of Model C at 1000 training samples mark, would require around double the training size.

From the current implementation, it can be seen that Model B requires 1000 more training days (i.e. training samples) to match the performance of Model C at the training size of 1000 days. Therefore,

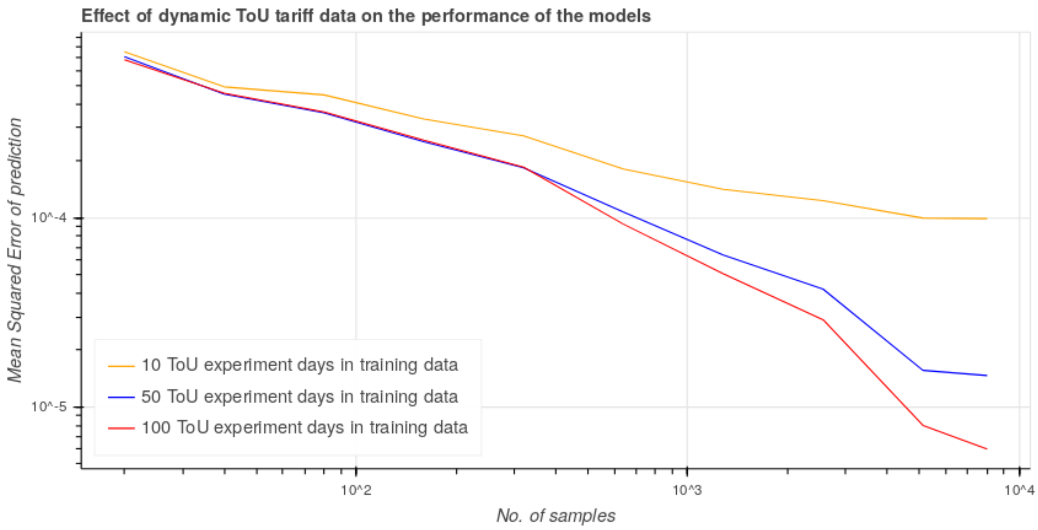

Figure 7.7: The performance comparison of the three models each being trained on the datasets which include 10, 50 and 100 ToU tariff days respectively. When the number of dynamic ToU tariff days is kept constant, we observe the effect of additional fixed tariff days

for the given case, the average effect of one ToU tariff experiment day is equal to 20 fixed tariff days. Therefore, in the ideal situation, the utility company should maintain the cost of one ToU tariff experiment day to be less than or equal to the cost of operation of 20 days of fixed tariff service to provide profitable ToU tariff schemes.

Of course, these results are only valid for the given model and the given number of training samples.

7.5. Summary

This chapter provides the first insights into the implementation phase of this research project. The chapter introduced the features and the feature selection process is explained. After selecting the features, a simply fixed tariff electricity demand forecasting model is considered and then the features importance plot is observed to verify the selection of the features. Then the problem setting is altered to predict the dynamic ToU tariff average electricity demand of users. We have seen the effect of fixed tariff days and ToU tariff days on the forecasting model. We have observed that increasing the number of fixed tariff samples to the ToU tariff demand forecasting model improves the performance due to the improvement of prediction for non-peak time blocks. In the next chapter, we will look into various active learning methods and the implementation of the two methods for selective sampling.

8
Active learning analysis

This chapter described the active learning methods used to selectively sample the ToU training samples for the forecasting model shown in chapter 7. The goal is to design an algorithm which improves the accuracy of a forecasting model with the limited budget of training ToU tariff experiment day. The active learning algorithm may pose queries to perform ToU experiment on a specific day. To improve the accuracy of the algorithm, it becomes essential to pick the 'most interesting' days, such that, information gained from the ToU tariff experiment will be most valuable. In this chapter, such algorithms are proposed and discussed. The chapter also discusses the machine learning framework used for each algorithm. At the end of the Chapter, the method of performance evaluation is discussed.

8.1. Random sampling vs selective sampling

The previous chapter considered the forecasting of energy consumption under dynamic ToU tariff policies. The selection of these days for the ToU tariff consumption was done randomly. Random selection of data may miss the important information hidden in the input subspace. Random selection of the data points for a training set may not result in the optimal solution of a machine learning model. This gives rise to the next important question; If we can only have limited number of ToU experiment days, at the same time want to improve the accuracy of the model, how can we achieve that? One way of achieving the goal is to carefully select the experiment days based on the potential information hidden in those days.

This leads us to investigate the active learning algorithms for achieving the goal using the algorithmic solution. Mind that, the forecasting model implemented in this chapter will be used to forecast the average energy consumption of the users; Active learning algorithms will only be used to find the days which will provide more information than other days and hence ToU tariff response for such days will be considered. The data gathered from this process will be used as training data for the forecasting model shown in the previous chapter.

8.2. General framework

As seen in chapter 3, there are three scenarios of active learning framework, namely - 1. Membership Query Synthesis, 2.Stream-based Selective Sampling, and 3.Pool-based Sampling. Out of these three frameworks, this research project uses a stream-based selective sampling method.

Let us consider the real-world scenario for the problem. The problem is best suited for online learning framework, where a stream of data is constantly flowing. Similar to the online learning method, every new energy consumption reading is generated at a particular interval (with 30 minutes interval). When a 'new' day is queried, the algorithm randomly selects the next day for the active learning analysis. The active learning algorithm then analyses the input feature space for the given tariff policy. If the algorithm decides to perform the ToU tariff experiment, then the ToU tariff along with the corresponding demand is considered for the sample day. Otherwise, the fixed tariff and corresponding fixed tariff demand are considered. The details about the data is explained in the chapter 5 and chapter 6.

Another new concept we will discuss in this chapter is the concept of informativeness $\mathcal{I}$. The active learning algorithm tries to quantify the informativeness of a data point with some measure; This measure is based on the method of active learning. The informativeness does not always directly indicate the information content of a data point, but it could also be an indicator of interestingness which indirectly provides knowledge of informativeness. For example, if a part of feature subspace is uncertain for the model, querying information about such input space would likely to be beneficial.

The overall flow of the active learning framework is explained below. The selective sampling for the Time-of-Use tariff experiments is a tricky concept as it is not just a single data point classification problem for a timeseries data, but it is a classification of a set of 48 data points. The informativeness $\mathcal{I}_i$ of every energy consumption observation should be considered while calculating the informativeness of a ToU experiment day. To recall the setting of ToU experiment setup, the energy consumption observation of users was collected with a half-hourly frequency. As mentioned earlier in this paragraph, the 'informativeness' of each energy consumption sample is quantified based on the analysis of input feature space. The average value of informativeness is then calculated by taking the mean of 48 values. In general,

$$\mathcal{I} = \frac{1}{48}\sum_{i=1}^{48} \mathcal{I}_i \tag{8.1}$$

The active learning algorithm queries the data from the data simulator presented in chapter 6. The data simulator randomly selects the next day sample and processes the demand response for the day. The active learning strategy then calculates the informativeness from the input space of the day sample and then compares the threshold $\mathcal{T}$. Then, based on the decision, it considers ToU tariff energy response or fixed tariff energy response, respectively. For informativeness I, the decision about the ToU tariff day is taken by using the following rule,

$$\text{Selective sampling criteria}: \begin{cases} \text{if } \mathcal{T} < |\mathcal{I}|, & \text{then, perform ToU experiment day} \\ \text{if } \mathcal{T} > |\mathcal{I}|, & \text{then, fixed tariff day} \end{cases}$$

The threshold $\mathcal{T}$ is manually set and decided as per the required number of ToU tariff samples. Ideally, the active learning model is then re-trained after a specific number of such decisions to update the knowledge of the environment. But, for understanding the fundamental differences between various methods and input feature space, the model is not re-trained unless mentioned otherwise. The analogy for this type of implementation is like follows - if a person is trained to press a big red button to note some unknown interesting thing happening behind a glass window, how well he/she performs in detecting valuable events. We can, therefore, assess the person's ability based on the quality of training given by training instructor (e.g. quality of training set) or the total time they spent in training (e.g. a number of training samples) or how does their mind decide what is interesting (e.g. type of active learning strategy).

Following pseudo-code shows the basic framework of active learning algorithm:

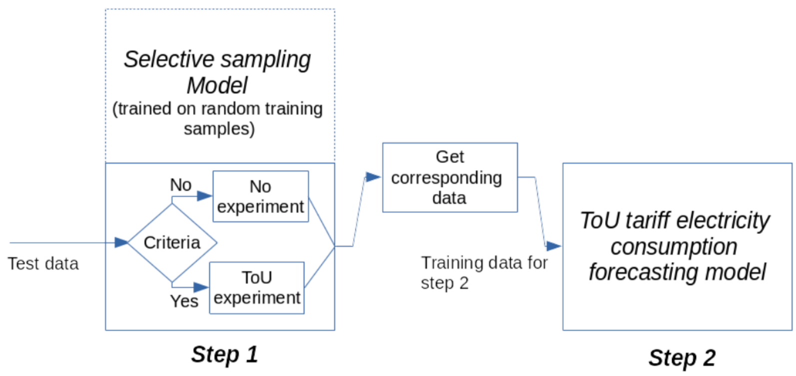

Figure 8.1: Active learning model design

Algorithm 2 Active learning basic algorithm

procedure Active learning(target variable, input features)
  Get n initial random ToU tariff consumption samples
  Get corresponding m number of input features
  Train active learning algorithm on the n number of ToU tariff consumption samples
  for i<total_days **do**
    Select test data for a day
    Calculate the informativeness l value from the given input space of the day
    if $\mathcal{I}>\mathcal{T}$ **then**
      Consider the ToU response consumption under ToU tariff policy as a target variable
    else
      Consider fixed tariff energy consumption data as a target output variable
      Change the dynamic tariff policy from input space to a fixed tariff policy
    Save the electricity consumption, tariff and other weather and calendar features in a dataset
  Train the forecasting algorithm on the generated dataset
  Deploy the model (in our case, use test data) to analyse the performance

As we can see from the algorithm 2, we are using two models to implement the active learning scheme. See figure 8.1 for reference. The selective sampling algorithm is trained on the randomly selected ToU tariff days. This model is then used to calculate the informativeness of each training day of the second model. The calculated informativeness is compared with the threshold $\mathcal{T}$ and decision about ToU tariff experiment is taken. According to the decision, corresponding data is stored in the training dataset of the next model. The next model (which is an actual ToU tariff forecasting model) is trained on that data. The performance of ToU tariff electricity consumption forecasting algorithm depends on the quality of selection by selective sampling algorithm.

8.3. Active learning strategies

Following sections go through the implementation of the active learning strategies considered to study under this research project.

8.3.1. Variance reduction

The first algorithm developed during the thesis is a novel adaptation of the QBC approach to incorporate the regression type target observations of a complete day (48 data samples). The variance reduction approach of active learning for regression setting is shown in section 3.3.3. In this method, the active learning model queries about the instance which it considers hard to predict (uncertain about single-point prediction). As seen in chapter 3, there are multiple ways to quantify the uncertainty. The focused method considers the variance of predictions provided by the group of machine learning regression models. Let x_i be the input feature set of $1 \times n$ size which include the features mentioned in section 7.4 of chapter 7. For a given day, the machine learning model maps the input x_i to target output y_i and the function is given by,

$$\hat{y}(x_i) = f(x_i) \qquad \text{for i = 1,2,...,48} \tag{8.2}$$

Where, $f(x_i)$ is the machine learning model which takes half-hourly values of feature variables x_i as an input to get predicted target variable $\hat{y}(x_i)$ for the actual target observation $y(x_i)$. To get the variance of the output predictions, we need to create a committee of the machine learning models; We will use random forest model with p number of high variance decision trees which are trained on slightly different data using the process called bagging, as explained in chapter 3. We then calculate the variance of predictions using

$$\mathcal{I} = \frac{1}{48}\sum_{i=1}^{48} \text{Var } f_{1:p}(x_i) \tag{8.3}$$

Where, Var $(\hat{y}(x_i))$ is the prediction variance. But this is the prediction variance for a single data point; We have 48 data points in a day. As our active learning strategy has to select a complete day, we need to consider the prediction variance for a complete day. This requirement leads us to the following equation which quantifies the informativeness of a day,

$$\mathcal{J} = \frac{1}{48}\sum_{i=1}^{48} \text{Var } f_{1:p}(x_i) \tag{8.4}$$

(8.5)

The data for a random day is selected from the dataset. The informativeness for each day is calculated using equation 8.5. Then the decision about considering the ToU tariff experiment is taken by comparing the informativeness $\mathcal{J}$ with a threshold; This threshold is a predefined constant value. If $\mathcal{J}$ is greater than the threshold then we will consider the energy consumption data related to the ToU tariff for that day. The entire process of variance reduction algorithm is shown here, in the Algorithm 3.

Algorithm 3 Uncertainty Sampling: Variance Threshold Method

procedure Variance Threshold($D(y_{\text{tou, q}}, y_{\text{fixed, q}}, X_p)$)
  Get n initial random ToU tariff consumption samples $y_{\text{tou},n}$ of size $m \times 1$ (where $n < q$)
  Get corresponding m number of features denoted by X_n of size $n \times m$
  Create a dataset $d_n(y_{\text{tou},n}, X_n)$
  Train the committee of p active learning models on $D(y_{\text{tou},n}, X_n)$ as shown in equation 8.2
  Set $count = 1$.
  for count < total_days do:
    Query the next day sample $d_1(y_{\text{tou},1}, y_{\text{fixed},1}, X_1)$ from $D(y_{\text{tou, q}}, y_{\text{fixed, q}}, X_p)$, where dimensions of X_1 are $48 \times m$
    Get prediction variance $\mathcal{J}$ from input space X_1
    if $\mathcal{J} > \mathcal{T}$ then
      Choose the ToU response consumption under ToU tariff policy as a target variable $y_{\text{tou},1}$
    else
      Choose fixed tariff energy consumption data $y_{\text{fixed},1}$ as a target output variable
      Change the dynamic tariff policy from input space to a fixed tariff policy
    Add the target variable and input data to the output dataframe D_{final}
    Increment the $count$ by 1 and repeat.
  Train the ToU tariff consumption forecasting algorithm on D_{final}
  Deploy the model (in our case, use test data) to analyse the performance

The data $D(y_{(\text{tou, fixed})}, X)$ is generated by the data simulator which will be discussed in chapter 6. A day-wise data (the 48 observations of target variables and $48 \times m$ values of input features) is queried from the above dataset for the inspection under the active learning algorithm. Once the data is queried, it is removed from the original dataset $D(y_{(\text{tou, fixed})}, X)$, making sure no sample is repeated.

8.3.2. Novelty detection

The hypothesis for the active learning algorithm based on novelty detection theory is that if a model can identify rare (but potentially informative) events, then the model performance may improve for certain days. In 1990, one of the biggest surges in sporting history still goes to the World Cup Semi-Final in 1990 - England v West Germany. The electricity demand was 2,800MW - equivalent to 1,120,000 kettles being boiled. Such events are rare in reality but have a high impact on the operations of the grid.

Similarly, the month of August we usually get sunny weather in the Netherlands. If we observe one heavy rain day during that month, it will be an interesting event to analyse the event as these events are rare but still have a non-zero chance of occurrence. Providing the machine learning model information about the feature subspace which was previously undiscovered may help us to improve the overall performance of the model.

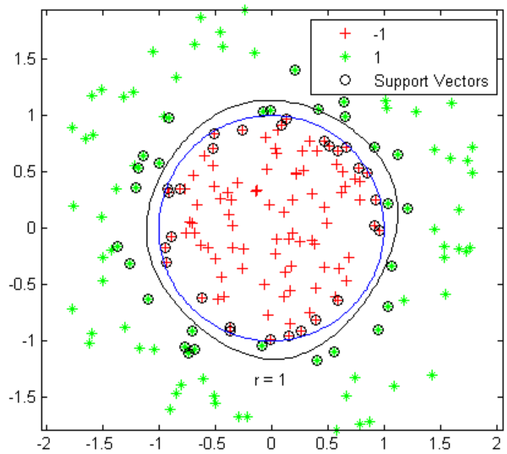

Figure 8.2: Example of one class SVM classification for 2D feature space

The next active learning algorithm is a one-class classification algorithm. Figure 8.2 shows the classification strategy for this method. The one class classifier is an unsupervised classifier which draws a boundary between the inliers and outliers. From the literature point of view, this algorithm fits more in the anomaly detection or novelty detection domain. The purpose of one-class classification is to find the novel data in a stream of data or to find an anomaly in the pool of data. This research project tries to analyze the effect of performing ToU tariff experiments on novel data points (days with unique input conditions) on the performance of the forecasting model. We will use one-class SVM to find the novelty in the input feature data. SVM with 'rbf' kernel is used to implement the one-class classifier.

The SVM algorithm is trained on the random selection of ToU tariff response days. The trained model forms a boundary line on the hyperplane of feature space. The decision function of an SVM algorithm provides the signed distance to the separating hyperplane. Where the signed distance is positive for an inlier and negative for an outlier. By considering the decision function of the SVM algorithm, we can vary the number of selected novel data points by dynamically selecting the threshold. The rest of the algorithm is similar to the variance reduction algorithm shown in the last section

Algorithm 4 shows the novelty detection algorithm.

This algorithm will make sure that the selected days for ToU tariff experiments are unique in the input space. Therefore, we are expecting to see interesting results from this algorithm.

The following section will cover the results of both the methods and analysis of the results will be presented.

8.4. Results

As shown in figure 8.1, we will use two model system for the active learning analysis. The purpose of the first model is to selectively sample the days for ToU tariff experiments and the second model is a ToU tariff consumption forecasting model, which is trained on the data collected by the selection criteria of the first model. See algorithms 2, 3 and 4 for more information.

The performance plots are generated using the Random Forest regression model which forecasts the ToU tariff consumption of users. The tariff policy is binary in the nature i.e. it only includes 'HIGH' and 'NORMAL' signals. Therefore, for a non-experiment day (i.e. a day with fixed tariff), only 'NORMAL' signal is used to indicate the fixed tariff signal.

The following section will talk about the results obtained by the two selected methods - (i) Variance reduction (ii) Novelty detection. First, we will see the compare the performance of our active learning model with the two baseline models. Then we will consider the effect of the number of ToU tariff samples (in the mixed training data of fixed and ToU tariff days) on the performance of active learning forecasting model. Then the effect of training size for selective sampling algorithm is analyzed. Finally, both variance reduction and novelty detection algorithm performance are compared with random selection algorithm and the conclusions are drawn.

Algorithm 4 Uncertainty Sampling: Variance Threshold Method

procedure Variance Threshold($D(y_{\text{tou, q}}, y_{\text{fixed, q}}, X_p)$)
    Get n initial random ToU tariff consumption samples $y_{\text{tou},n}$ of size $m \times 1$ (where $n < q$)
    Get corresponding m number of features denoted by X_n of size $n \times m$
    Create a dataset $d_n(y_{\text{tou},n}, X_n)$
    Train the one class SVM model on $D(Y_{\text{tou},n}, X_{1:n})$ to generate a boundary on the SVM feature hyperplane.
    Set $count = 1$
    for count < total_days do
        Query the next day sample $d_1(y_{\text{tou},1}, y_{\text{fixed},1}, X_1)$ from $D(y_{\text{tou, q}}, y_{\text{fixed, q}}, X_p)$, where dimensions of X_1 are $48 \times m$
        Get the average distance l of the input data X_1 from the decision boundary.
        if $l > \mathcal{T}$ then
            Consider the ToU response consumption under ToU tariff policy as a target variable $y_{\text{tou},1}$
        else
            Consider fixed tariff energy consumption data $y_{\text{fixed},1}$ as a target output variable
            Change the dynamic tariff policy from input space to a fixed tariff policy
        Add the target and input data to the output data frame D_{final}
        Increment the $count$ by 1 and repeat.
    Train the ToU tariff consumption forecasting algorithm on D_{final}
    Deploy the model (in our case, use test data) to analyse the performance

8.4.1. Comparison with baseline models

In this section, we will study the effect of selectively sampled data on the performance of the forecasting model. We will consider various scenarios to investigate the performance of the forecasting model.

Experiment setting

For this experiment, we consider three models-

- (baseline 1): The model is trained ONLY on the ToU tariff observations
- (baseline 2): The model is trained ONLY on the fixed tariff observations
- Active learning model: The model is trained on selectively sampled ToU days (by variance reduction method)

All three models of this experiment use the same base data for training. Out of the three models, Baseline 1 model is trained on ToU tariff data of the training days. Baseline 2 models are trained on fixed tariff data of the training days. The active learning model uses a combination of ToU and fixed tariff days. The process is explained below.

The selective sampling (variance reduction) algorithm of the active learning model is trained on 500 randomly selected ToU tariff days. This model is used to classify the training data of ToU tariff consumption forecaster of the active learning model using the variance of predictions by the committee. The threshold for classification is set at $\mathcal{T} = 0.0004$.

The three ToU tariff consumption models with the same hyperparameters are trained on the following number of training days - [20, 40, 80, 160, 320, 640, 1280, 2560, 5120, 8000]. The models are then tested on a common test set of 500 ToU tariff days. The MSE of the predictions is plotted against the number of training samples. The above experiment is run 5 times and the average of the results is plotted in figure 8.3.

Discussion: Overall results

The active learning model has classified 443 out of 8000 days for ToU tariff experimentation. Therefore, this data includes 443 ToU tariff days and 7,557 fixed tariff days. The performance in figure 8.3 suggests that the active learning algorithm performed better than the baseline 1 algorithm, as expected. In theory, the baseline 1 model should perform poorly as it has the least knowledge of the response consumption of the energy users. Also, a supervised model should have superior performance as it has

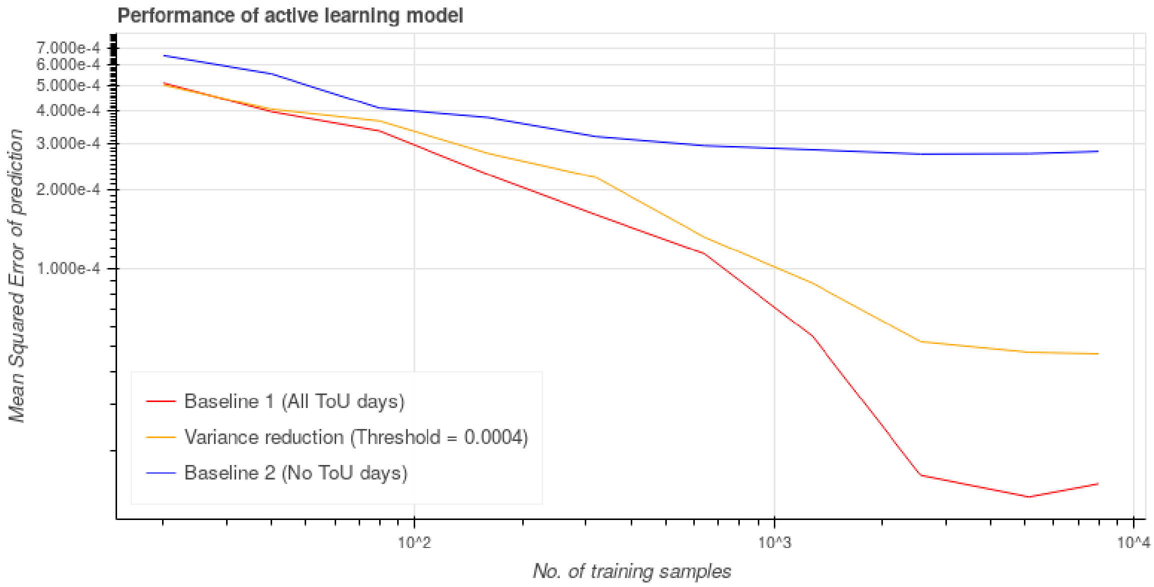

Figure 8.3: The above three models are trained on three types of data - Baseline 1: All ToU days, Baseline 2: No ToU days and Active learning (variance reduction) : mixed days (ToU and Fixed). The active learning model has performed 443 total experiments in 8000 days of data. The mean squared error is the considered evaluation measure for all three algorithms

the most knowledge about the response consumption for ToU tariff signals. The similar outcome can be observed from the figure 8.3.

Discussion: Comparison of active learning performance with Baseline 1 performance

Now we will compare Baseline 1 and active learning (variance reduction) method by their mean squared error plots. The active learning classified only 443 days for ToU tariff experiment. Therefore, the active learning model is only introduced to 443 different ToU tariff policies. Even with so few numbers of ToU experiments, we see that the active learning model improves the predictions with the help of fixed tariff training days. Also, we see that the performance of both models is very similar for a small number of training samples as both models include a high proportion of ToU tariff days in the training data. After the first 443 ToU tariff days of training data, when the remaining training data are supplied with fixed tariff days, we see a slight decline in the performance of the active learning model. The Baseline 1 model improves significantly over a large number of training days and saturated at a lower MSE value. It is also worth noting that the active learning model's performance for a large number of training data (443 ToU days + other fixed tariff days) is a significant improvement over the Baseline 2 performance.

Discussion: Effect of ToU tariff signals on Baseline 2

It is an interesting detail of the above experiment. If the Baseline 1 model has never experienced HIGH tariff, how will it respond to it? or would it respond at all? It should be considered that a tariff signal is a categorical value and not continuous value. So, the mathematical relationship between a tariff and output is non-linear. We know that the random forest model is based on decision trees, which essentially create rule-based nodes. The various features will have various importance level in the trees. For example, we know that mostly the largest residential loads are thermal loads. We also know that temperature has a high influence on thermal load usage. Therefore, temperature as a feature should have high importance in the above random forest model (and it has!). But the fixed tariff signal has no information about electricity consumption. It will have the same fixed value for the lowest and highest energy consumption points. Interestingly, the tariff signal has the least importance for Baseline 1 model as shown in figure 8.4. The change in the value of tariff signal will have zero effect on the Baseline 1 model because the value of tariff signal does not have any effect on the Baseline 1 model.

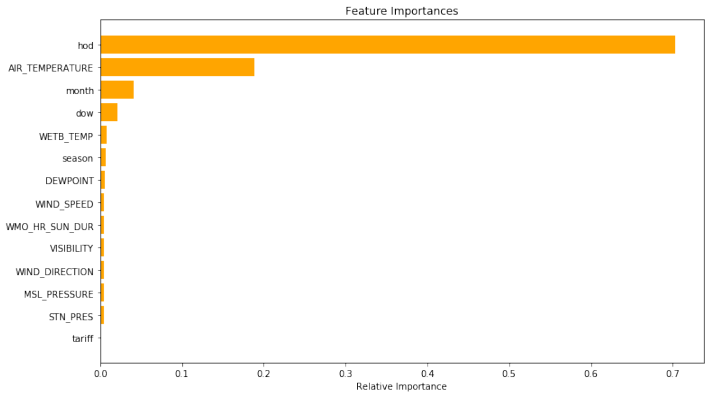

Figure 8.4: Feature importance chart for Baseline 2 model (No ToU tariff samples)

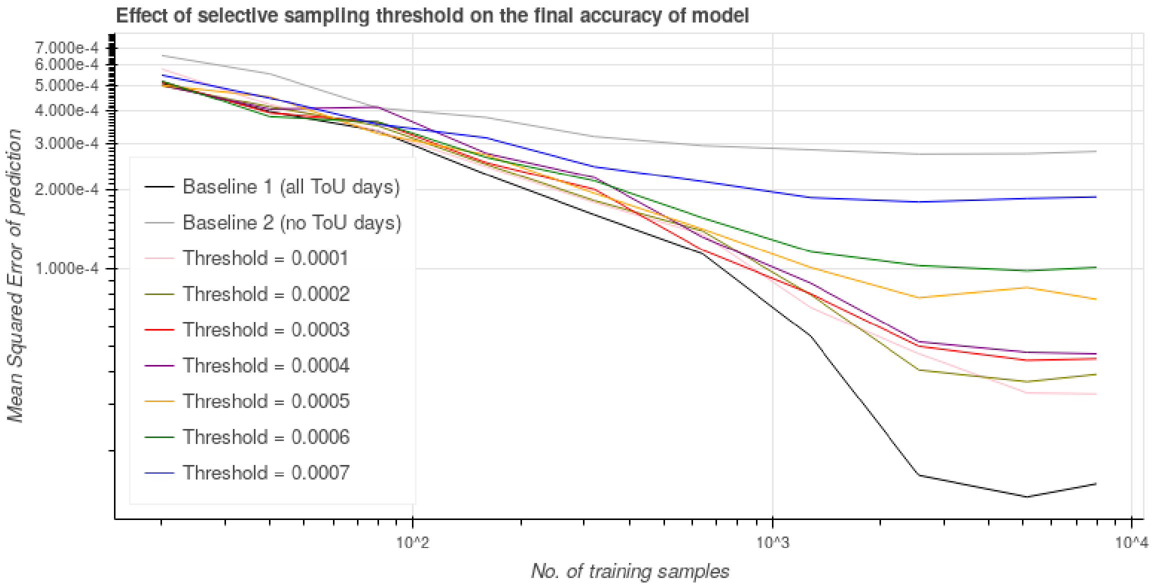

Figure 8.5: Effect of threshold on the performance of the active learning model

8.4.2. Effect of threshold on the performance of active learning

Next, we study the effect of threshold on the number of days which qualify for ToU experiments. The data sampled by different values of the threshold is used to train forecasting models. The performance of these models is then plotted.

Experiment setting

The selective sampling model is trained on 500 randomly selected ToU tariff days. 7 different values of the threshold are considered. With each threshold value, the training data for the ToU forecasting model is gathered as shown in figure 8.1. Then 7 ToU consumption forecasting models are run to analyse the performance of models with each threshold level. The above experiment is run 5 times and the results are averaged. Similar to the previous experiment setting, two baselines are considered - i) no ToU days ii) all ToU days. MSE performances of these models for an increasing number of training sets are plotted. The three ToU tariff consumption models with the same hyperparameters are trained on the following number of training days - [20, 40, 80, 160, 320, 640, 1280, 2560, 5120, 8000]

Discussion: Threshold vs the number of qualified days

Table 8.1 shows the effect of threshold value on the number of ToU tariff experiment queries issues by the algorithm. It is no wonder that as the threshold is decreased, more samples become qualified for the ToU tariff experiment query. Figure 8.5 graphically shows the effect of the threshold on the

Threshold	0.0001	0.0002	0.0003	0.0004	0.0005	0.0006	0.0007
Queries per 1000 samples	478	231	119	51	17	7	3
Total queries	3822	1852	951	409	135	56	25

Table 8.1: Effect of active learning threshold value on a number of queries and performance of the algorithm. The total queries are calculated for the sample size of 8,000 days.

performance of the algorithm. As the lower threshold qualifies more samples for ToU experiments, we see that lower threshold results in lower MSE, thereby, increasing the accuracy of the model.

Discussion: Performance of models

The comparison between the performance of Baseline 2 and the performance of the algorithm (with threshold = 0.0007) provides insights into the power of a few ToU tariff samples. The later algorithm only issues an average of 3 queries per 1000 samples. Whereas Baseline 2 issues no queries at all, leading to complete blind spot for ToU tariff conditions. Just with few queries, active learning algorithm, or to more generalize, the semi-supervised algorithm can perform much better than a model which is not exposed to ToU experiment conditions.

8.4.3. Trade-off between training size of selective sampling method and queries

The selective sampling (variance reduction) algorithm calculates the prediction variance of each input day. This variance is dependent on the prediction confidence of a model. In this section, we see the effect of the size of the training set on the output prediction variance of the selective sampling model.

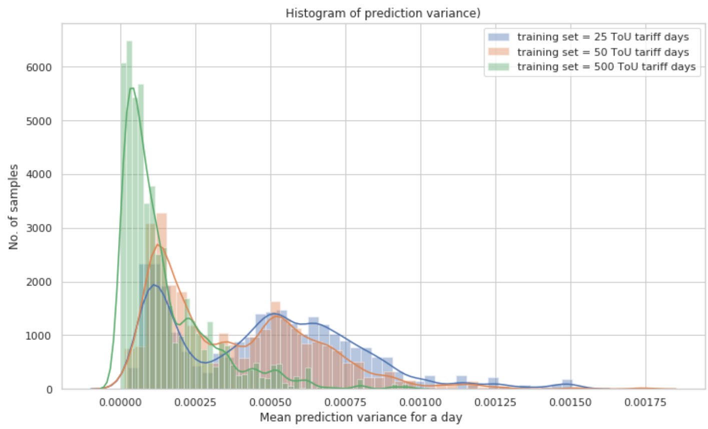

Figure 8.6: Effect of number of training samples on prediction variance of selective sampling model(variance reduction). The three selective sampling models are trained on following number of randomly selected ToU tariff days - 25, 50, 500

Experiment setting

(In the later phase of the thesis, it was found that the ToU tariff consumption forecaster algorithm used during the implementation of this experiment had a faulty implementation. That implementation randomly selected the training set and testing set from the same pool of data. Therefore, it was possible to have the same data in training and testing dataset. Therefore, the results and discussion related to the faulty implementation are omitted in the and attached in the Appendix. Following discussion contains the results from the verified implementation)

The following experiment is carried out with three selective sampling models (variance reduction) with a different number of random training samples. First selective sampling model (TS1) is trained on only 25 randomly selected ToU days. The next model (TS2) is trained on 50 randomly selected ToU days. The third model (TS3) is trained on 500 randomly selected ToU days. Then each model is deployed to classify the 8,000 days into ToU experiment day and non-ToU experiment day. The results are plotted on a histogram.

	Threshold	Training samples	Queries per 1000 samples	Total Queries
TS 1	5E-4	25	348	2781
TS 2	5E-4	50	276	2210
TS 3	5E-4	500	17	141

Table 8.2: (Variance reduction model) Effect of a number of training samples on the number of queries and accuracy of the models. The total queries are calculated for the sample size of 8,000 days.

Discussion: training set size vs number of ToU queries for the selective sampling algorithm

Figure 8.6 shows the histogram for the effect of a number of training samples for selective sampling (variance reduction) model on the variance of prediction (a measure of informativeness $\mathcal{I}$). As we increase the number of training samples for the selective sampling (variance reduction) model, the confidence of model increases, resulting in lower prediction variance. The selective sampling model trained on large training set will have higher confidence in the predictions, thus will classify a smaller number of data points for ToU tariff experiments. Here, we see two different trade-offs- 1) between the training samples and query samples 2) between quality and quantity of queries. If we try to cut the cost of data labelling for training in an earlier stage, it will reflect in the higher cost of data labelling for each query produces by the model.

To understand the effect more clearly, table 8.2 provides more insights about the effect of a number of training samples on the number of sample queries and the final prediction error of the model. We see that when the selective sampling model is trained on a higher number of training samples, the queries generated by that model are low. When we increase the number of instances for query decision, we see that TS3 queries much fewer labels than TS2 and TS3.

In the next section, we will analyse the random sampling and selective sampling criteria and will observe the effect of training samples on the performance of the model.

8.5. Results: passive vs active learning

As the research project revolves around the selection of the data points for the ToU tariff demand 'labelling', it is important to finally look at the performance of active vs passive 'label' selection process. In this section, we will compare the results of passive learning and active learning.

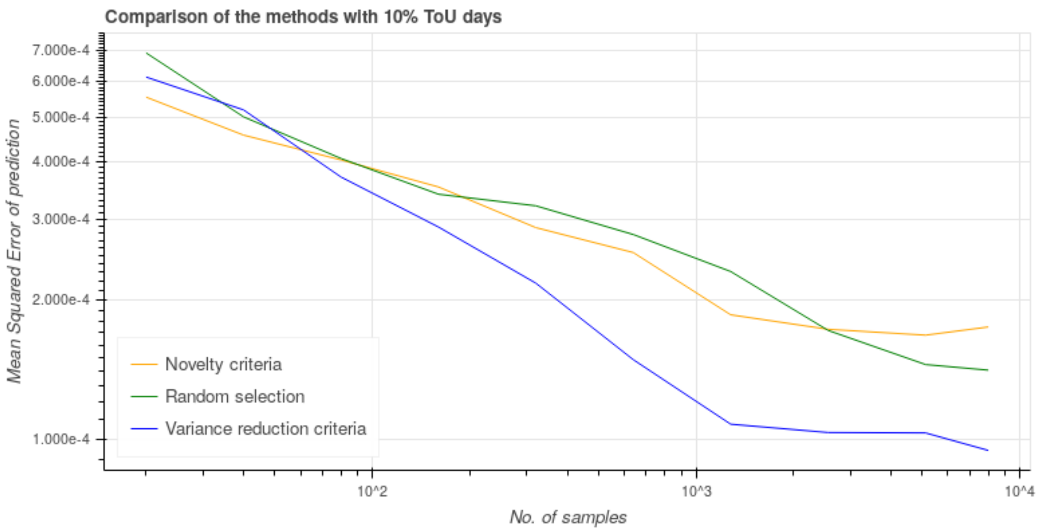

Figure 8.7: Performance of active and passive learning models

Experiment setup

The selective sampling (variance reduction and novelty detection) algorithms are trained on 500 days of the average ToU tariff energy consumption. Then these algorithms are tested on the input feature set of 8000 days. The threshold is set such that only top 10% days (800 days) will be qualified for the

ToU tariff experiment days (for variance reduction model: threshold = 0.00035; for novelty detection model:$\nu = 0.1$, $\gamma = 0.1$). Similarly, a third dataset is made from those 8000 samples by randomly selecting 800 days for the ToU tariff experiment observations. The three ToU tariff consumption models with the same hyperparameters are trained on the following number of training days - [20, 40, 80, 160, 320, 640, 1280, 2560, 5120, 8000]. The 10% of each training set is reserved for ToU tariff days. For example, a training set of 20 days will contain 18 randomly selected fixed tariff days and 2 ToU tariff days (selected either by active or passive selection criteria) and so on. Then these models are tested on the 500 days of ToU tariff input features for predicting the ToU tariff demand observation. The above procedure is repeated 5 times and the average of the MSE of each round is plotted against the size of the training set.

The figure 8.7 shows the performance of the three forecasting model whose training data is generated by various active and passive techniques explained above.

Discussion: overall results

For the smaller size of training sets, all three models show similar performance (with novelty detection algorithm performing slightly better). But, as the number of samples increases, we see a steep decrease in the MSE of variance reduction method, whereas, novelty sampling and random sampling show a lower degree of decrements in the MSE of the prediction. The novelty detection method achieves the saturation at much higher MSE value than random selection and variance reduction methods. Finally, we see that the variance reduction algorithm performs best out of the three methods. Whereas, the performance of novelty detection and random sampling is at a comparable level with the random sampling achieving lower MSE above the training set of 5000 days.

Discussion: novelty algorithm performance

With the smaller size of training sets, the novelty detection algorithm performs slightly better than the other two models. But when more training data is introduced to the models, the performance of the forecasting model based on novelty detection samples starts declining. As we increase the number of training samples, the novelty detection algorithm starts training its forecasting model on more number of novelties. This leads the model to have a lot of information about novelties also known as out of distribution data. The knowledge of this data may provide limited knowledge about the actual distribution of data. Therefore, the knowledge pool of the novelty detection model has a large gap regarding the 'normal' (in distribution) ToU day. More number of training samples with such information biases the novelty detection model to have a stronger belief on a certain set of knowledge which may not be true for most of the data. We can see that for a very large training set (above 2,000 days), the novelty detection algorithm performs worse than the random selection strategy. The random selection criteria make sure that for a large number of training samples, the selected ToU tariff days are (more or less) uniformly picked from the distribution of underlying data. Therefore, for larger training sets random selection method may have selected enough data points which can provide valuable information about the significant number of ToU experiment test days. Therefore, we see that the performance of the random selection model beats the performance of the novelty detection model.

8.5.1. Comparison of selective sampling by two methods

This section analyzes the data points which are selected by the two active learning algorithms. We will analyze the two selective sampling algorithm outputs by the visualization of data.

The analysis of selective sampling algorithm (variance reduction) of the active learning model is performed to check relation about prediction variance between the trees of random forest algorithm. The selective sampling algorithm is trained on 500 randomly selected ToU tariff days. And the test data of 500 separate ToU tariff days are used to calculate the prediction variance using the above model. As we have actual observations of ToU tariff days, we can calculate the mean squared error of the predictions. The scatter plot between the prediction variance and actual prediction MSE is shown in figure 8.8. We see that as prediction MSE increases, the prediction variance also increases. It should be noted that the relationship is not perfectly linear and the data points with very high prediction error may not show high prediction variance.

Now, we compare the data points selected under the two strategies. Figure 8.9 shows the comparison plot between the data selections by the novelty detection and variance reduction algorithms. The scatter plots of prediction variance and MSE of predictions is shown in the figure. The electricity consumption

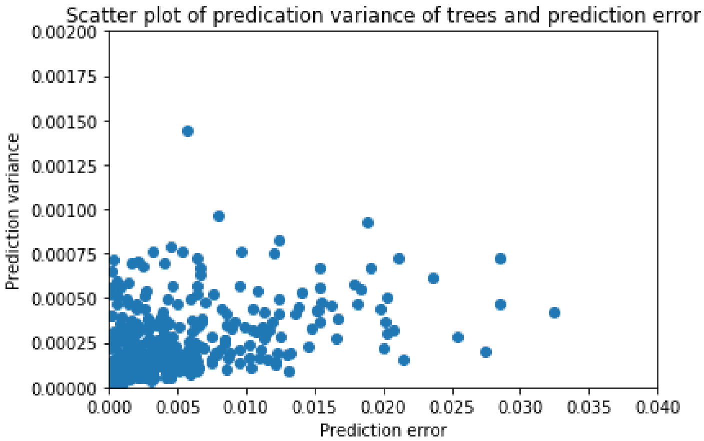

Figure 8.8: Prediction variance vs prediction error scatter plot

for 500 days is considered for this selective sampling analysis. Now we see the effect of the threshold level for the variance reduction and novelty detection criteria; For example, the 10% selective sampling criteria are applied with novelty detection and variance reduction criteria, where each algorithm selects 50 most important days from the pool of 500 days according to their criteria.

Figure 8.9: The comparison of data points selected by variance reduction vs novelty detection from the same pool of 500 days. The subplots are the scatter plot of prediction variance vs MSE of predictions for the same pool of 500 days (data points). Blue colour data points are selected under the novelty detection algorithm, orange colour data points are the non-novelty data points. The size of the data point is directly proportional to the distance of the data point from the decision boundary of SVM classifier.

The variance reduction algorithm selectively samples all the data points above the threshold (dotted red line). The sampling by the novelty detection algorithm is shown by blue data points. The data points which are not selectively sampled are shown in orange colour. These comparison plots show us that both the algorithms sample a different set of days for ToU tariff experiments. Therefore, we can expect the ToU tariff forecasting model to learn different information from this training data and potentially have different results.

To investigate the selection criteria further, we will analyze the distribution plots of the features of the selectively sampled features. The same set of 500 days of observations is considered for the following plot. For this distribution plotting, mean values of temperature and dew point observations of every day are considered; Whereas, the categorical variables like a day of week or month does not need any alterations. Figure 8.10 shows the distribution plots of the following features - temperature, dew point, day-of-week, month. The variance reduction and novelty detection algorithms are tested on the same

500 days of data points shown in figure 8.9. The models are trained on randomly selected, unbiased 500 days from the pool of data generated by the data simulator from chapter 6.

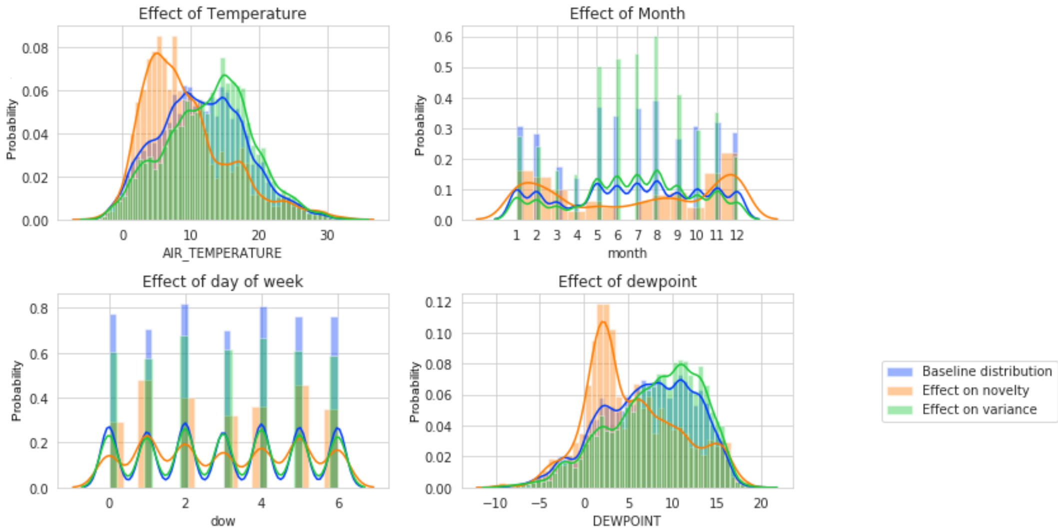

Figure 8.10: Distribution of features of selectively sampled data points is observed for the same 500 data points analysed in figure 8.9. Investigated features are (clockwise from top left) - temperature, month, dew point, day-of-week

It is interesting to see that the selective sampling by the variance reduction algorithm has the distribution which is similar to the original distribution of the testing data. Whereas, the selectively sampled data from the novelty detection algorithm has a distribution which is sharp and does not follow the underlying distribution of test data. We see more winter data points are selected than any other season. Perhaps, the selective sampling method focuses on the data with highly variable characteristics i.e. winter. For example, the electricity consumption pattern is highly variable throughout the winter day. This variation in electricity consumption might be categorised as a 'novel' situation. In any case, it is interesting to see the focus of the novelty detection algorithm on the data points with highly variable electricity consumption patterns.

8.6. Summary

In this chapter, the implementation of the active learning algorithm is discussed. The active learning framework is explained for variance reduction and novelty detection algorithms. Then the performance of the active learning algorithm is benchmarked with the supervised and unsupervised algorithms in similar settings. The effect of training samples of the active learning model on the final forecasting performance is evaluated. Finally, we compared the active learning algorithms with the passive learning algorithm (where a day is randomly selected for ToU tariff experiment observations). We have seen that the variance reduction algorithm performs better than the random selection criteria, but novelty detection shows comparable performance for the given setting.

9

Conclusion and future work

In this chapter, the conclusions of the research project are drawn. The conclusions and potential improvements in the form of future work are presented in the following chapter. First, the primary conclusions are discussed which mainly revolve around the performance of active learning models. Then the secondary conclusions are presented.

9.1. Conclusions

The main aim of the research project was to investigate active learning strategies for ToU tariff experiment day selection. The problem was not a traditional active learning classification problem. The problem setting was an active learning timeseries regression, the domain which has seen limited research[37]. Therefore, to tackle the unique nature of the problem, a novel QBC based algorithm was developed by taking the inspiration from some of the work done in the regression domain with variance reduction criteria[31]. One class classifier was another algorithm which was investigated to actively select the ToU tariff experiment days.

The research problem has a distinct setting. Usually, in the active learning setting, classification problems are solved by finding the most informative data points which are usually around the decision boundary. Other than the queried data points, no other data point has labels. In our case, we have two types of data points - ToU tariff consumption and fixed tariff consumption. We aim to forecast electricity consumption under ToU tariff. In this setting, we do not have truly unlabelled data, as, the fixed tariff consumption and ToU tariff consumption of the same day share about 75% of electricity consumption observations (only the consumption at peak tariff prices is different).

This setting makes the research problem a difficult case for analysing active learning techniques. Chapter 7 has shown that a model trained on 30% randomly selected ToU tariff days can perform quite close to the model trained on all ToU tariff days. This result has shown that the introduction of any random ToU tariff day to the model will improve the model performances significantly. Also, in the case of the model with 30% randomly selected training samples, we see the fixed tariff training days improve the accuracy of the model as we train the model on a higher number of samples. This effect is seen clearly in figure 8.5, where a model trained with just 25 ToU tariff days (with rest of the fixed tariff training days) improves the performance significantly and outperforms the model trained on fixed tariff day by an adequate margin.

The same figure 8.5 shows a comparative study of multiple models with a different proportion of ToU tariff samples in a training set. We see that the introduction of more ToU days results in better performance. We need more analysis of the cost of an experiment to comment more about the performance of additional experiment days.

The analysis of quality of samples vs quantity of samples has shown that the selective sampling algorithm of active learning model trained on a small number of ToU training days will query more days for the ToU tariff experiment as the variance of the ensemble model is large. In fact, we have seen that the selective sampling model TS1 which was trained on randomly sampled 500 ToU tariff days, queried 17 days per 1000 samples, which is much lower than other two query sizes: TS2 - 25 training set, 348 queries per 1,000 samples and TS3: 50 training set, 276 queries per 1,000 samples. When we increase the number of training sample of a selective sampling model, we increase the confidence of the model by generating strong causal relations between the input and target variable of the data. Therefore, an ensemble model trained on a larger set of training data will lead to higher confidence in the prediction of test data, lowering the prediction variance between the committee of the decision trees. The histogram of the prediction variance shown in figure8.6 supports the above hypothesis.

We see that variance reduction algorithm performs better than the novelty detection algorithm. To analyse the results, we plot the distribution of queries along with the distribution of base data in figure 8.10. When we compare the distribution of queries by the algorithms, it is observed that the distribution of variance reduction algorithm is in similar to the distribution of the original data. Whereas, the distribution of novelty detection algorithm tries to capture the data which has a lower probability of occurrence. It is also likely that the novelty detection algorithm queries the data points which have high variation i.e. during the winter season the electricity consumption of residential users vary much more than the other seasons.

This leads to a discussion about the selection criteria of the two methods. The variance reduction method tries to query the data points which create high disagreement between the committee of decision trees. Usually, this disagreement is inversely proportional to the confidence about the knowledge of that feature subspace. Whereas, the novelty detection algorithm queries the unique data points. In other words, the algorithm picks the data points which do not fit in the distribution of the base data. Now, as we have discussed the underlying selective sampling philosophy of these methods, we try to formulate a possible conclusion about the performance of three methods -

1. Variance reduction: The variance reduction method performs well when compared to the novelty detection algorithm and random selection. Now, this result can mean one of two possibilities (or the combination of both) - either the model is focusing on the less known feature space by querying more data from low probability distribution area of the input space, or the model is trying to learn the input data which has more complex relationship with the target variable. We observe that the distribution of the queries by the variance reduction model is within the distribution of the base data (figure 8.6). Therefore, it is concluded that the variance reduction model queries the data point with input variables which have complex relationships with the target variable (electricity consumption). The model does not focus the queries explicitly on unique or new data i.e. out-of-distribution data.

2. Novelty detection: Unlike variance reduction method, the novelty detection algorithm explicitly focuses the queries on the unique data. The poor performance of the novelty detection algorithm can be traced to the figure 8.10 which shows the distribution of the input space of queries sampled by novelty detection and variance reduction. This plot also provides the distribution of a base dataset. We can see that novelty detection algorithm samples more days from low probability region, as the low probability feature space qualifies for 'novelty'. But, by training the active learning model explicitly on the out-of-distribution days, the novelty detection model creates a knowledge gap. If we increase the number of selectively sampled (by novelty detection) ToU tariff days in the training data of training set of a forecasting model, the model builds high confidence about the relation between input and output variables of the training data. This does not guarantee the performance improvement of the model. The knowledge gain by this kind of selection method (based on 'uniqueness') may not be completely useful as the test data may not follow the same distribution. The test set is randomly selected from the pool of data. Therefore, the test data follows the distribution of actual data and has the most data points which follow the distribution of the actual data. As the distribution of test data is different than the distribution of training data, we see comparatively poor results for novelty detection algorithm.

   It is worth noting that even though the performance of the novelty detection algorithm is poorer than the variance reduction algorithm, the performance of the novelty detection algorithm improves by increasing the size of training data. This leads to another conclusion that the out-of-distribution data also shares some concepts with the test set which is uniformly distributed over the input dataset.

3. The random selection strategy shows mixed results for the implemented case. The random selection strategy works well for a lower number of training samples as well as a very high number of training samples. First, we will focus on the performance of the model on a lower number of

training data. For training size of 20, 40 and 80 sample days, we get only 2, 4 and 8 ToU tariff days respectively (10% of total data). For such a small number of ToU days, the queries generated by active learning algorithm may not provide any significant information about the test data. As we increase the number of training data, more number of ToU days qualify in the training set. As the number of training samples increases, we see that the performance of the variance reduction model improves much quicker than random selection model.

For a high number of training samples (larger than 2000), we see that the rate of improvement for variance reduction method has slowed down, but random selection model has relatively better rate improvement. In fact for the training set of 8,000 days with 10% ToU tariff days, the random sampling method outperforms the novelty selection strategy by a significant margin. For a very large training set with 10% ToU tariff days, random selection criteria samples ToU days (more or less) uniformly from the distribution of input data. Therefore, forecaster based on randomly selected ToU days may have a better knowledge of test data feature space than the ToU days selected by the novelty detection algorithm.

9.1.1. Discussion on secondary outcomes of the research

This section focuses on the secondary outcomes of the thesis. These outcomes are discussed chapter by chapter.

Chapter 6: Demand Response Simulator

The main outcome of the demand response simulator is the generation of ToU tariff policy and corresponding electricity consumption of users. The motivation behind designing the load switching latency was to reflect the complex user behaviour. In chapter 7, we observe the performance of the fixed tariff consumption forecasting model and ToU tariff consumption forecasting model. The MSE curve of fixed tariff consumption does not reach to noise bed after 8,000 training days, indicating that the model can learn more about the consumption pattern by increasing the training samples. Whereas, the MSE performance curve of ToU tariff consumption forecasting model reaches noise bed around 2,000 training samples. This indicates that the complexity of the data can not be fully learned by the model. The reader should mind that both the models are initialized with the same hyperparameters. The above discussion provides validation for the design choices of demand response simulator.

Chapter 7: Basic forecasting and analysis

Firstly, we will talk about the results of feature importance plot for fixed tariff consumption forecaster from chapter 7. It is observed that the hour-of-day and air temperature are the top features for the fixed tariff consumption forecaster. We have already shown the causality between temperature pattern and electricity consumption pattern in chapter 5.

Chapter 8: Active learning analysis

Let us discuss the dilemma of optimizing the number of labels for the active learning based forecasting model. In the section 8.4.3 we have seen that, for the same threshold, lower training samples for active learning algorithm results in a higher number of queries, whereas, a higher number of training samples for active learning model will result in a lower number of queries. This optimization problem is very hard to solve in a generalized manner. The subjective analysis of each problem setting will help provide the optimal number of total labels.

9.2. Further work

Following are the suggestions for the future work:

- Due to lack of the data, a data simulator had to be developed during the implementation phase of this project. Therefore, the results are limited by the design of the demand response simulator. More reliable results can be obtained if real-world data for the ToU tariff policies are adequately available for training a forecasting model with acceptable performance.
- This project considers a limited number of features, due to time and resource constraint of the master's thesis. More research can be done on the feature identification for ToU tariff consumption forecasting. Also, as the weather data had to be gathered from another source, the exact location of the weather station and the residential users may not have matched.

- Instead of the binary tariff policy considered in the thesis, a more complex ToU tariff policy can be implemented to improve the knowledge about the effect of the tariff policy on active learning performance.

- The active learning models of this thesis are designed in two steps- The first step will be trained on randomly selected ToU days to classify the test data into ToU and non-ToU experiment days (refer figure 8.1). The classified data is used to train the ToU tariff consumption forecaster while the original training data of the selective sampling model is discarded.

  The above design was implemented to analyse the effect of queries on the final forecast of the active learning models. This implementation can be improved to incorporate the initial training data of selective sampling model into the training data for the forecasting model of the same active learning algorithm. This implementation has the potential to change the results of the novelty detection algorithm as currently the algorithm only considers out-of-distribution data for training of the forecasting model.

- In the current implementation, the selective sampling model of active learning model is only trained once at the start. By retraining the selective sampling and forecaster models of active learning algorithm on the previous test data, we can improve the knowledge gain with the same number of total queries. This iterative active learning model may potentially have a higher convergence rate.

- Lastly, an investigation about the effect of the accuracy of selective sampling model on the final accuracy of forecaster can be done. This approach can provide further information about the importance of query quality

A

Appendix A

(The following results are obtained from the faulty implementation of the ToU tariff forecasting model. The implementation allowed the training data to be a part of the testing data. Therefore, the results of the implementation are not reliable. This mistake was found in the later phase of the thesis. Due to lack of time, the MSE performance plot could not be regenerated. Therefore, the following discussion related to section 8.4.3 is attached in the appendix.)

The following experiment is carried out with three selective sampling models (variance reduction) with a different number of random training samples. First selective sampling model (TS1) is trained on only 25 randomly selected ToU days. The next model (TS2) is trained on 50 randomly selected ToU days. The third model (TS3) is trained on 500 randomly selected ToU days. Then each model is deployed to classify the 8,000 days into ToU experiment day and non-ToU experiment day. The results are plotted in a histogram.

The three ToU tariff consumption models with the same hyperparameters are trained on the following number of training days - [20, 40, 80, 160, 320, 640, 1280, 2560, 5120, 8000]. This experiment is repeated 10 times and the average of the MSE values are plotted in the performance plot.

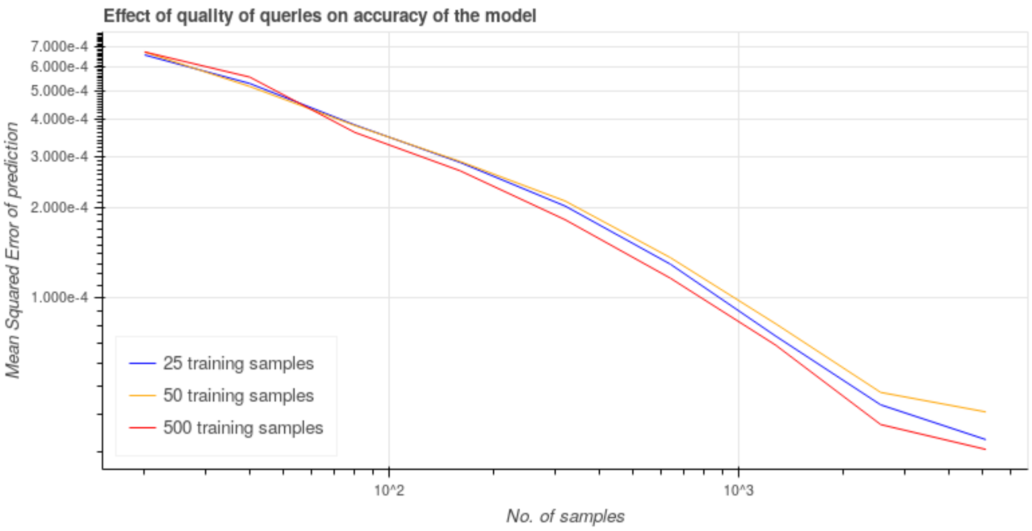

Figure A.1: Effect of number of pre-training samples on final accuracy of the model

The effect of the number of initial training samples of the variance reduction active learning model on the final accuracy of the ToU tariff demand forecasting model is graphically shown in figure A.1. This figure showcases the results of three different active learning models (with a different number of training samples). The x-axis shows the number of training samples for all three models. These training samples are a mix of fixed tariff days and ToU tariff days. As shown in figure 8.6, as the active learning model is trained on the different number of training samples, the number of queries formed by these models will vary for the three cases.

Figure A.1 shows the performance of all three models with MSE as a performance measure. The models are trained on the mentioned number of ToU training days and other fixed tariff days. For example, 100 training samples of TS1 will include 25 ToU tariff days and 75 fixed tariff days. The models are tested on the test dataset of 500 ToU tariff days. The experiment is repeated 10 times with randomised dataset and the mean of the results has been reported to reduce the variance in the results. Interestingly, the active learning model trained on 25 ToU tariff samples (TS 1) performs better than a model which is trained on 50 ToU tariff samples (TS 2). But how is it possible to get better results with a model which is less informative about the feature space? The answer to this interesting fact lies in the active learning theory and findings of the table 8.2. An active learning model creates queries which represent the uncertainty of the model. The final machine learning model is trained on a collection of queries created by the active learning model. Therefore, as seen in table 8.2, TS 1 creates the highest number of queries followed by TS 2 and then by TS3. But as discussed at the starting of this section, it is not just the quantity of data which matters, but the quality of data also matters. Even if TS 2 creates more number of queries than TS 3, that does not necessarily mean that the quality of the queries is equal for both the models. Theoretically, TS 3 would produce higher quality queries as it is trained on a much higher number of training samples. With this new and better understanding of the setups, we can better understand figure A.1. For a lower number of samples, model based on TS 3 data performs worse than the other two models. But as we increase the number of samples, we get to see a clear trend that model based on TS 3 data performs the best, followed by TS 1 and TS 2. Here, we see the trade-off between quantity and quality of data. It is easy to realise the cost associated with the quantity of ToU tariff samples, but it should be noted that TS 3 samples with higher quality also have hidden cost. We can not neglect the initial ToU tariff samples used for training of TS 3.

Bibliography

[1] Mygridgb. URL http://www.MyGridGB.co.uk/.

[2] Jamshid Aghaei and Mohammad-Iman Alizadeh. Demand response in smart electricity grids equipped with renewable energy sources: A review. Renewable and Sustainable Energy Reviews, 18:64–72, 2013.

[3] Mohamed H Albadi and Ehab F El-Saadany. A summary of demand response in electricity markets. Electric power systems research, 78(11):1989–1996, 2008.

[4] Dana Angluin. Queries and concept learning. Machine learning, 2(4):319–342, 1988.

[5] Jacob Benesty, Jingdong Chen, Yiteng Huang, and Israel Cohen. Pearson correlation coefficient. In Noise reduction in speech processing, pages 1–4. Springer, 2009.

[6] Margaux Brégère, Pierre Gaillard, Yannig Goude, and Gilles Stoltz. Target tracking for contextual bandits: Application to demand side management. arXiv preprint arXiv:1901.09532, 2019.

[7] David Cohn, Les Atlas, and Richard Ladner. Improving generalization with active learning. Machine learning, 15(2):201–221, 1994.

[8] David A Cohn, Zoubin Ghahramani, and Michael I Jordan. Active learning with statistical models. Journal of artificial intelligence research, 4:129–145, 1996.

[9] John J Conti, Paul D Holtberg, Joseph A Beamon, A Michael Schaal, JC Ayoub, and James T Turnure. Annual energy outlook 2014. US Energy Information Administration, page 2, 2014.

[10] Soma Shekara Sreenadh Reddy Depuru, Lingfeng Wang, Vijay Devabhaktuni, and Nikhil Gudi. Smart meters for power grid—challenges, issues, advantages and status. In 2011 IEEE/PES Power Systems Conference and Exposition, pages 1–7. IEEE, 2011.

[11] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. The elements of statistical learning, volume 1. Springer series in statistics New York, 2001.

[12] Samuel Gyamfi, Susan Krumdieck, and Tania Urmee. Residential peak electricity demand response—highlights of some behavioural issues. Renewable and Sustainable Energy Reviews, 25:71–77, 2013.

[13] Rabab Hassan and Ghadir Radman. Survey on smart grid. In Proceedings of the IEEE SoutheastCon 2010 (SoutheastCon), pages 210–213. IEEE, 2010.

[14] David F Hendry and Kirstin Hubrich. Combining disaggregate forecasts or combining disaggregate information to forecast an aggregate. Journal of business & economic statistics, 29(2):216–227, 2011.

[15] Sheng-Jun Huang, Rong Jin, and Zhi-Hua Zhou. Active learning by querying informative and representative examples. In Advances in neural information processing systems, pages 892–900, 2010.

[16] Masahito Ito, Ryuya Uda, Satoshi Ichimura, Kazuya Tago, Tohru Hoshi, and Yutaka Matsushita. A method of appliance detection based on features of power waveform. In 2004 International Symposium on Applications and the Internet. Proceedings., pages 291–294. IEEE, 2004.

[17] Jeffrey C Jackson. An efficient membership-query algorithm for learning dnf with respect to the uniform distribution. Journal of Computer and System Sciences, 55(3):414–440, 1997.

[18] Ajay J Joshi, Fatih Porikli, and Nikolaos Papanikolopoulos. Multi-class active learning for image classification. In 2009 IEEE Conference on Computer Vision and Pattern Recognition, pages 2372–2379. IEEE, 2009.

[19] Irena Koprinska, Mashud Rana, and Vassilios G Agelidis. Correlation and instance based feature selection for electricity load forecasting. Knowledge-Based Systems, 82:29–40, 2015.

[20] Anders Krogh and Jesper Vedelsby. Neural network ensembles, cross validation, and active learning. In Advances in neural information processing systems, pages 231–238, 1995.

[21] Joseph C Lam. Climatic and economic influences on residential electricity consumption. Energy Conversion and Management, 39(7):623–629, 1998.

[22] David D Lewis and William A Gale. A sequential algorithm for training text classifiers. In SIGIR'94, pages 3–12. Springer, 1994.

[23] Andy Liaw, Matthew Wiener, et al. Classification and regression by randomforest. R news, 2(3): 18–22, 2002.

[24] London Project Low, Carbon. Low carbon london - project closedown report. pages 1–19, 2005.

[25] Shuai Lu, Nader Samaan, Ruisheng Diao, Marcelo Elizondo, Chunlian Jin, Ebony Mayhorn, Yu Zhang, and Harold Kirkham. Centralized and decentralized control for demand response. In ISGT 2011, pages 1–8. Ieee, 2011.

[26] Naoki Abe Hiroshi Mamitsuka et al. Query learning strategies using boosting and bagging. In Machine learning: proceedings of the fifteenth international conference (ICML'98), volume 1. Morgan Kaufmann Pub, 1998.

[27] Tom M Mitchell and Machine Learning. Mcgraw-hill science. Engineering/Math, 1:27, 1997.

[28] Easton Nick. A day in the life of an operational energy manager. National Grid Transco, 2010.

[29] Kamal Nigam, Andrew McCallum, Sebastian Thrun, Tom Mitchell, et al. Learning to classify text from labeled and unlabeled documents. AAAI/IAAI, 792:6, 1998.

[30] J. Ross Quinlan. Simplifying decision trees. International journal of man-machine studies, 27(3): 221–234, 1987.

[31] Tirthankar RayChaudhuri and Leonard GC Hamey. Minimisation of data collection by active learning. In Proceedings of ICNN'95-International Conference on Neural Networks, volume 3, pages 1338–1341. IEEE, 1995.

[32] Neil Rubens, Mehdi Elahi, Masashi Sugiyama, and Dain Kaplan. Active learning in recommender systems. In Recommender systems handbook, pages 809–846. Springer, 2015.

[33] Arthur L Samuel. Some studies in machine learning using the game of checkers. ii—recent progress. In Computer Games I, pages 366–400. Springer, 1988.

[34] James Schofield. Dynamic time-of-use electricity pricing for residential demand response: Design and analysis of the low carbon london smart-metering trial. 2015.

[35] Bernhard Schölkopf, Robert C Williamson, Alex J Smola, John Shawe-Taylor, and John C Platt. Support vector method for novelty detection. In Advances in neural information processing systems, pages 582–588, 2000.

[36] Burr Settles. Active learning literature survey. Technical report, University of Wisconsin-Madison Department of Computer Sciences, 2009.

[37] Burr Settles and Mark Craven. An analysis of active learning strategies for sequence labeling tasks. In Proceedings of the conference on empirical methods in natural language processing, pages 1070–1079. Association for Computational Linguistics, 2008.

[38] H Sebastian Seung, Manfred Opper, and Haim Sompolinsky. Query by committee. In Proceedings of the fifth annual workshop on Computational learning theory, pages 287–294. ACM, 1992.

[39] Claude Elwood Shannon. A mathematical theory of communication. Bell system technical journal, 27(3):379–423, 1948.

[40] Jasmina Smailović, Miha Grčar, Nada Lavrač, and Martin Žnidaršič. Stream-based active learning for sentiment analysis in the financial domain. Information sciences, 285:181–203, 2014.

[41] Ingo Stadler. Power grid balancing of energy systems with high renewable energy penetration by demand response. Utilities Policy, 16(2):90–98, 2008.

[42] Johan AK Suykens and Joos Vandewalle. Least squares support vector machine classifiers. Neural processing letters, 9(3):293–300, 1999.

[43] James W Taylor and Roberto Buizza. Using weather ensemble predictions in electricity demand forecasting. International Journal of Forecasting, 19(1):57–70, 2003.

[44] Cynthia A Thompson, Mary Elaine Califf, and Raymond J Mooney. Active learning for natural language parsing and information extraction. In ICML, pages 406–414. Citeseer, 1999.

[45] S Tindemans, P Djapic, J Schofield, T Ustinova, and G Strbac. Resilience performance of smart distribution networks. 2014.

[46] Simon Tong and Daphne Koller. Support vector machine active learning with applications to text classification. Journal of machine learning research, 2(Nov):45–66, 2001.

[47] Leslie G Valiant. A theory of the learnable. In Proceedings of the sixteenth annual ACM symposium on Theory of computing, pages 436–445. ACM, 1984.

[48] John S Vardakas, Nizar Zorba, and Christos V Verikoukis. A survey on demand response programs in smart grids: Pricing methods and optimization algorithms. IEEE Communications Surveys & Tutorials, 17(1):152–178, 2014.

[49] Meng Wang and Xian-Sheng Hua. Active learning in multimedia annotation and retrieval: A survey. ACM Transactions on Intelligent Systems and Technology (TIST), 2(2):10, 2011.

[50] Xingquan Zhu, Peng Zhang, Xiaodong Lin, and Yong Shi. Active learning from data streams. In Seventh IEEE International Conference on Data Mining (ICDM 2007), pages 757–762. IEEE, 2007.

[51] Indrė Žliobaitė, Albert Bifet, Bernhard Pfahringer, and Geoff Holmes. Active learning with evolving streaming data. In Joint European Conference on Machine Learning and Knowledge Discovery in Databases, pages 597–612. Springer, 2011.

[52] Ahmed Zoha, Alexander Gluhak, Muhammad Imran, and Sutharshan Rajasegarar. Non-intrusive load monitoring approaches for disaggregated energy sensing: A survey. Sensors, 12(12):16838–16866, 2012.