Department of Mechanical and Aerospace Engineering

# Project Title: Improving Energy Efficiency and Cost-effectiveness through Demand Forecasting of DHW Consumption in Residential Buildings: A case study of Findhorn, North Scotland

Author: Chilufya Chitindi

Supervisor: Dr Paul Gerard Tuohy

A thesis submitted in partial fulfilment for the requirement of a degree in

Master of Science in

*Sustainable Engineering: Renewable Energy Systems and the Environment*

August 2023

*Student No. 202290074*

# Copyright Declaration

# Abstract

The global drive towards sustainable energy sources with acceptable green credentials has led to a bias towards the electrification of various sectors, including transportation. This includes the heating and cooling in business, residential and commercial buildings. Some communities are establishing decentralised energy systems or microgrids to incorporate more renewable energy sources. To manage the challenges of these distributed energy resources on the broader grid, flexibility markets are emerging. These markets optimise local energy resources and reduce costs, by utilising advancements in digitisation and artificial intelligence. These have given rise to the ubiquitous smart grids, enhancing energy generation, distribution, and consumption efficiency. The Findhorn ecovillage, in Moray, Scotland, while dedicated to sustainability, grapples with significant energy inputs and leveraging technology could address these challenges. This study analyses the hourly hot water demand of the end user and utilises a machine learning algorithm, an ARIMA model developed in Python, to guide the development of a smart control system for charging a residential building's thermal store. The ARIMA models developed showed promising results, with low RMSE and MAE values in comparison to other models. The results of this study offer valuable insights that can be leveraged in future initiatives to optimise energy consumption in Findhorn's residential buildings.

Keywords: DHW usage, Smart Grids, Time Series Analysis, Machine Learning, ARIMA.

# Acknowledgements

I would like to thank my supervisor Dr Paul Gerard Tuohy for his guidance and support throughout this project.

Additional thanks go to all the people who have encouraged and supported me throughout this challenge both personally and professionally, including but not limited to my family and friends, Dr Christopher Chileshe, for his unwavering support and kindness and most importantly to God, without whom this opportunity would not have been possible.

# Table of Contents

# List of Figures

# List of Tables

# Nomenclature

| Symbol | Description | Units |
|---|---|---|
| ACF | Autocorrelation Function | |
| ADF | Augmented Dickey-Fuller | |
| AIC | Akaike Information Criterion | |
| AI | Artificial Intelligence | |
| ANN | Artificial Neural Network | |
| ARIMA | Autoregressive Integrated Moving Average | |
| BIC | Bayesian Information Criterion | |
| DER | Distributed Energy Resource | |
| DHW | Domestic Hot Water | |
| DNO | Distribution Network Operator | |
| DSO | Distribution System Operator | |
| MAE | Mean Absolute Error | |
| MSE | Mean Squared Error | |
| OEM | Open Energy Monitor | |
| PACF | Partial Autocorrelation Function | |
| PV | PV: Photovoltaic | |
| RF | Random Forest | |
| RMSE | Root Mean Squared Error | |
| SMA | Simple Moving Average | |
| STL | Seasonal decomposition of time series | |
| SVR | Support Vector Regression | |
| SVM | Support Vector Machine | |

## 1.0 Introduction

The transition to renewable energy sources is a crucial step in the realisation of the global campaign popularly known as net zero, to reduce greenhouse gas emissions.[1]. Projections published by the former Department for Business, Energy, and Industrial Strategy (BEIS, now the Department for Energy Security and Net Zero - DESNZ) suggest continued growth in energy generation from renewable technologies in the UK through 2040 (Figure 1) [2].

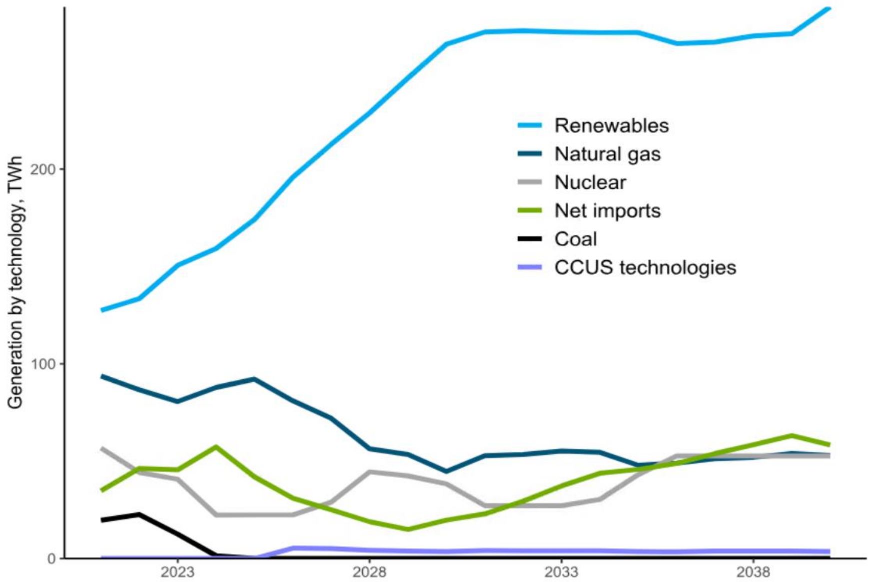


*Figure 1: Electricity generation by fuel source, TWh*

### 1.1 Smart Electrification

The growth of wind and solar technologies has significantly increased the supply of clean electricity in recent years. However, the demand side, primarily transportation and heating sectors, still mostly depend on fossil fuels [3]. Continued cost reductions in wind and solar technologies, have now made, renewable power a viable solution for electrifying many industries, both directly and indirectly. Strategic electrification offers an economical way to decarbonise these sectors, improve system flexibility, and incorporate larger proportions of renewable energy into power systems [4].

In order to attain the shared objective of a carbon-neutral future, the worldwide shift in energy utilisation is concerned not just with the methods of energy production but, crucially, with the patterns of its consumption as well. Essentially both the supply and demand aspects must be concurrently addressed to ensure the comprehensive and robust decarbonisation of the entire system. Consequently, innovations within end-use sectors are pivotal in facilitating this transition. The International Renewable Energy Arena (IRENA) suggests routes for the smart electrification of these industry sectors. These routes are illustrated in Figure 2 [4].

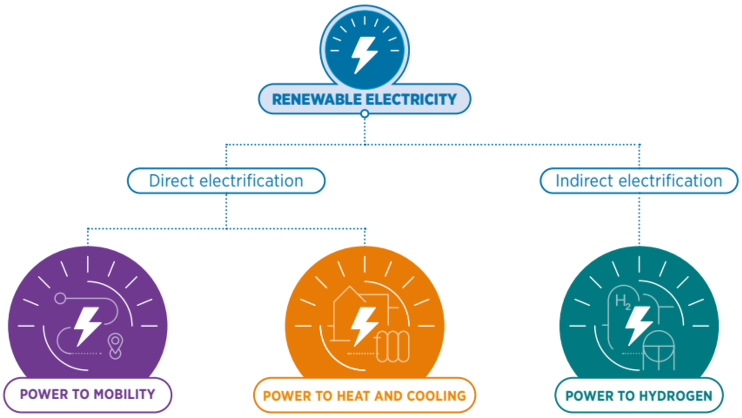


*Figure 2: Routes for smart electrification, source IRENA 2023 [4]*

## 1.2 Electrification of Heating and Cooling

Heating and cooling of buildings, which constitute about 50% of global final energy use, exceed the consumption of electricity (20%) and transportation (30%) [4]. These sectors also contribute to over 40% of the world's energy-related CO2 emissions [5]. The buildings and industrial sectors make up nearly 95% of global heating demand [7]. Predominantly, heat is generated from the use of fossil fuels and some unsustainable biomass sources. In recent times, the fossil fuel share in heat production has been declining, from 91% in 2010 to 75% in 2021 [6]. To achieve the 1.5°C target by 2050 and decarbonize these sectors, electrification, especially of heating, is key. Forecasts based on IRENA's 1.5°C Scenario suggest that electricity would fuel 73% of the total demand in buildings by 2050, a significant rise from the current 34% [7]. This transition would necessitate the deployment of about 793 million heat pump units by 2050, a 1267% increase from the current 58 million units [4].

## 1.3 Centralised and Decentralised Energy Systems

Centralised energy systems rely on large-scale power plants, offering efficiency and unified standards but posing vulnerabilities like single failure points and potential monopolistic practices. Decentralised systems, utilising sources like solar or wind, increase resilience, adapt to local needs, and often harness sustainable sources. However, they face grid integration challenges and can have higher initial costs [8].

An increasing number of local communities in the United Kingdom are forming decentralised energy systems or micro grids, which consist of small-scale renewable energy generation and local energy storage systems, that can offer a more sustainable and resilient alternative to centralised energy systems. Additional potential benefits of decentralised energy systems include reduced carbon emissions, increased energy efficiency and cost effectiveness, and greater energy security.

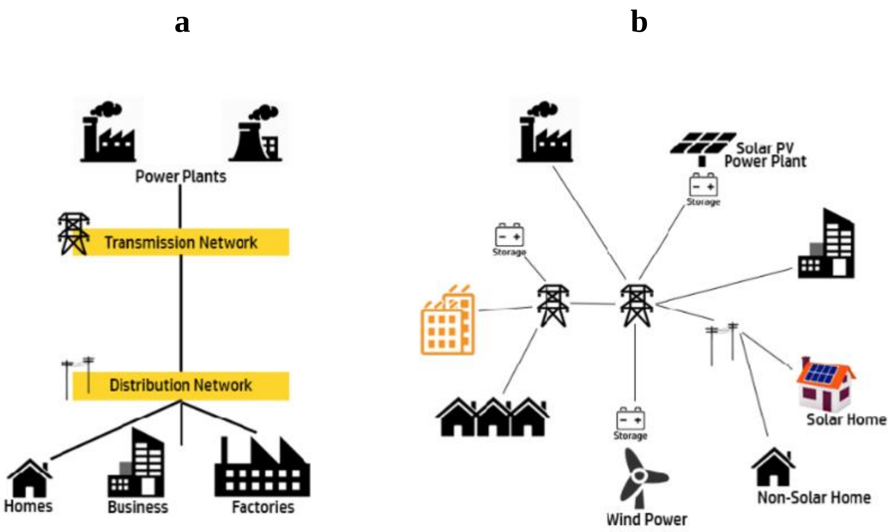


*Figure 3: a) Centralised vs b) Decentralised Energy Systems*

Despite these benefits, the decentralised nature of energy production can lead to a fragmented grid system, making it more challenging to ensure a stable and reliable energy supply. Without proper coordination and planning, this could result in inefficiencies and higher costs for consumers. It is therefore essential to address these challenges to ensure equitable access and a reliable energy system for all customers.

### 1.3.1 Distributed Energy Resources (DERs)

Distributed energy resources (DERs) refer to a diverse range of small-scale power generation and storage technologies that are located close to the point of consumption. These resources include solar photovoltaic (PV) systems, wind turbines, microgrids, energy storage systems, and electric vehicles. DERs play a crucial role in the development of local energy systems by providing decentralised and sustainable sources of energy

DERs offer several benefits such as the enhancement of the resilience of the energy system by reducing reliance on centralised power plants and transmission infrastructure. In the event of a grid failure or natural disaster, DERs can continue to operate and provide power to critical loads, thus improving the overall reliability of the system. They also enable local energy optimisation and demand response capabilities through advanced monitoring and control systems.

### 1.3.2 End-user Engagement in Flexibility Markets

Despite the benefits of DERs, distribution systems continue to encounter operational challenges from the unpredictability of renewable generation and DERs. Distribution system Operators (DSOs) are turning to market tools, like local flexibility markets, to manage these challenges by trading flexibility with participants, such as end-users and aggregators (an entity that consolidates energy assets, bridging asset owners with the flexibility market), for economic efficiency.

Flexibility Markets provide end-users with greater control and flexibility in managing their energy resources by becoming active participants in the energy market. Such engagement is made possible by the integration of technologies, such as smart grids and digital platforms, to enable user involvement and facilitate the exchange of energy and flexibility services.

However, the true potential of these decentralised energy systems can only be fully harnessed when comprehensive information is readily accessible to the market participants. Knowledge about when energy is needed and the optimal times to utilise available power are important. This insight ensures the success of flexibility markets, benefiting both end-users and system operators.

## 1.4 Enabling flexibility through Digitalisation

Digitalisation facilitates the development of strategies that leverage data from metering devices through computational programs or machine learning algorithms. These strategies provide flexibility, optimise operations, and reduce costs [9]. Metering devices are used to monitor energy use and operations enable new smart system models that help cut down on energy bills. Digitalisation therefore allows consumers to reduce energy costs and enables the electricity grid to optimise the use of renewable sources [10] [11].

## 1.5 Artificial intelligence and Data Analysis for forecasting

AI improves energy system flexibility by accurately predicting and optimising heating and cooling demands, facilitating greater use of fluctuating renewable energy sources. AI techniques, such as deep learning and time series analysis, lead to significant energy savings, particularly as they can detect when buildings are empty (which is more than 60% of the time) and can adjust control settings accordingly [12] [13]. In a 2019 Report, Artificial Intelligence and Big Data, the IRENA suggest steps to leverage Big Data and AI in the Power Sector (Figure 4), most of which were adopted in the current study.

AI enhances the functionality of IoT by adding intelligence to the data and communication between its components. It shows great promise in creating predictive energy consumption models for buildings, considering factors like the building's thermal properties, architecture, weather conditions, wind speed and direction, indoor and outdoor temperatures, and user behaviour. Recent studies indicate that AI-based methods improve prediction accuracy, enabling better forecasts of short-term changes essential for control applications [14].

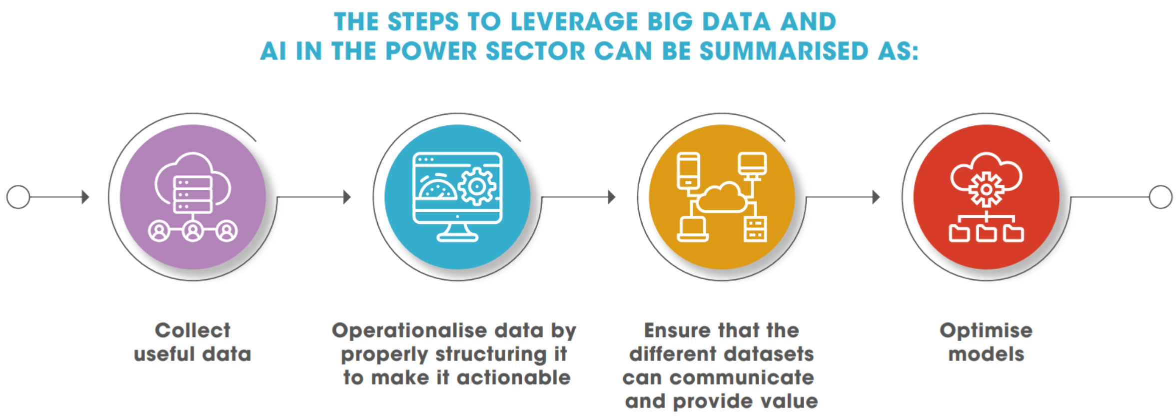


*Figure 4: Steps to Leverage Big Data and AI in the Power Sector [15]*

### 1.5.1 Artificial Neural Networks (ANN)

ANNs are a type of forecasting technique capable of learning from historical data and identifying patterns, enabling them to provide accurate predictions and optimise energy management strategies. ANNs can be trained to model the flexibility of smart grids by considering factors such as demand response, energy storage, and renewable energy generation. This modelling approach can assist in designing more efficient and resilient energy systems that can adapt to changing conditions and optimise resource allocation. ANNs can be applied in energy management and smart grids because incorporating flexibility models has the potential to address the challenges posed by the transition to sustainable energy systems [16].

### 1.5.2 Support Vector Machine (SVM)

SVM is a machine learning technique that focuses on finding an optimal hyperplane to separate data into different classes. In the context of energy consumption forecasting, SVM models are trained to classify historical data into low, medium, or high demand categories. These models have been applied to various energy forecasting tasks, including short-term load prediction and electricity price forecasting. SVM models have the ability to handle high-dimensional data and the potential for incorporating additional features such as weather conditions, socioeconomic factors, and time of day.

### 1.5.3 AutoRegressive Integrated Moving Average (ARIMA)

ARIMA is a  method used in forecasting, particularly within the realm of artificial intelligence, for predicting future trends or predicting electricity consumption patterns based on historical data.

The models consists of three components: 'AutoRegression' which assesses the dependency between a current observation and a number of its preceding ones, 'Integrated' which uses the differencing of raw observations to transform the data into a stationary time series, and 'Moving Average' that models the error term of the process as a linear combination of previous error terms [17].

Through these mechanisms, ARIMA models can enable precise forecasting of future electricity demand. This accurate forecasting is integral to the efficient operation of smart grids, as it helps in effective load management, optimises grid operations, and facilitates the integration of renewable energy sources.

### 1.5.4 Limitations of AI forecasting methods

All the above methods utilise AI to various extents. However, some of the limitations associated with AI based forecasting techniques for energy consumption are that the performance of these models heavily depends on the quality and availability of input data. Lack of data or inaccurate data can lead to unreliable predictions. There is a need for proper data preprocessing and feature selection techniques to enhance the accuracy of forecasting models. Additionally, the selection of appropriate model parameters and optimisation algorithms is crucial for achieving optimal performance [18].

Despite these limitations, implementing these AI techniques enables distribution network operators (DNOs) to make data-driven decisions, improve system performance, and enhance the integration of renewable energy sources to the electricity grid. Overall, there is a significant role for AI in enabling Big Data services in distribution networks and but further research and development is required in this field [19].

# 2.0 Project Aims and Objectives

The aim of the current study was to contribute to the current DHW consumption demand profiling methods by developing a demand forecasting algorithm to aid in optimising thermal store charging thereby reducing energy costs associated with domestic hot water use at the West Whins site in Findhorn.

## 2.1 Objectives

The specific objectives for the project included:

- To investigate the current demand forecasting techniques for DHW consumption
- To develop a demand forecasting algorithm to predict hourly DHW consumption in residential buildings in Findhorn
- To critically evaluate and validate the developed algorithm
- To propose a method of implementing the tool at West Whins.

## 2.2 High-level Approach

The steps to accomplish the study's goals are outlined as follows:

| | |
|---|---|
| **Project Plan** | **Step 1**: Review current literature to understand leading methods to inform the choice of tools and case study support. |
| | **Step 2**: Establish the technical methodology for analysis |
| | **Step 3**: Perform a technical analysis using selected tools and methods |
| | **Step 4**: Discuss the findings, address limitations and provide recommendations for future work. |

# 3.0 Case Study Background – Findhorn Ecovillage

Eco-villages, like Findhorn in the north of Scotland, strive to reduce their environmental impact by using local renewable energy sources such as wind and solar, integrating with the electricity from the grid. This has the added benefit of lowering the cost of purchasing electricity from the grid [20].

Figure 5: Findhorn Ecovillage Housing with Local Energy Sources

## 3.1 West Whins Housing Complex

The energy infrastructure at the Findhorn Ecovillage incorporates an internal microgrid connected to the national grid through a single sub-station. The microgrid combines wind and Solar PV electricity generation. Within this context, the West Whins site in the Findhorn Ecovillage is a housing complex with 6 flats.

At the West Whins site, there is a central energy centre (EC) responsible for supplying energy to the housing complex. The EC consists of a 14kW air source heat pump (ASHP) providing both space heating and DHW supply. Additionally, there is a 550-litre hot water thermal store specifically designed for DHW supply, and charged directly by 6 solar thermal panels; each measuring 2.3m$^2$. Billing for hot water consumption is calculated based on a single heat meter assigned per flat, takes into account flow meter readings and assumes a nominal temperature difference of 43$^o$C determined by a flow temperature of 55$^o$C and a cold water input of 12$^o$C [21].

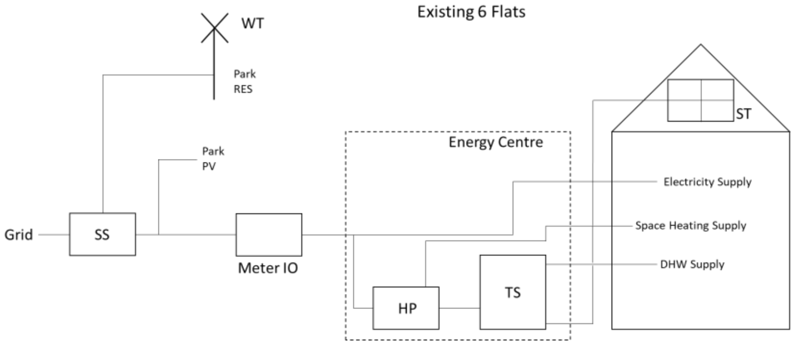


Figure 6: West Whins 6 Flat Structure

Figure 7: Smart Control Panels in the West Whins Energy Centre

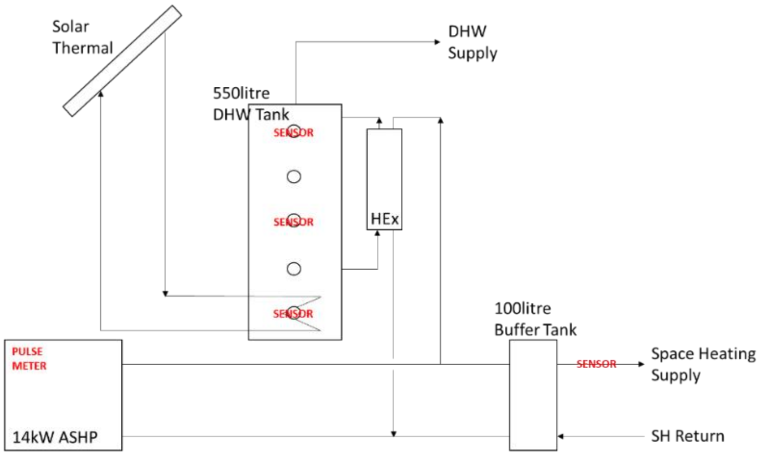


*Figure 8: a) DHW & Space Heating System at West Whins b) DHW Tank*

## 3.2 OpenEnergyMonitor Kit

The Energy and environmental monitoring and control function at Findhorn Ecovillage uses OpenEnergyMonitor (OEM) supply equipment. Particular choice of OEM equipment was largely influenced by its open-source nature, versatility in various applications, as well as the availability of case studies that align with the requirements at Findhorn.

EmonPi metering device is based on Raspberry Pi technology and is used to gather and upload data from the energy centre to a cloud based service. This device enables the connection of diverse sensors and provides the capability for continuous data uploading to the ‘Emoncms’ cloud service through Wi-Fi connectivity. Furthermore, the Emon Pi hub offers additional data storage capacity, enhancing the overall functionality of the monitoring system.

The device firmware primarily provides the energy monitoring function and handles the radio traffic from other sensor nodes. It utilises a software library named emonLibCM and is set up for two current channels to gather and process voltage and current measurements. Additionally, it receives temperature data from external sensors and manages pulse inputs. EmonCMS is the main user interface on the emonPi/base, it can be used to store and visualise data locally or just used to configure posting data to a remote server, or both.


*Figure 9: a) Emonpi metering device b) Emoncms user interface*

## 3.3 Previous Work

Knowledge of Findhorn's electrical network stemmed from the EU-funded FP7 Orchestration of Renewable Integrated Generation in Neighbourhoods (ORIGIN) Project, which sought to maximise local renewable usage and curtail energy imports in established buildings. Later research under the Smart Integrated Energy Systems (SIES) Project addressed newer buildings, aiming to reduce energy consumption for users and refine the park's microgrid. Yet, during the SIES project, the West Whins systems could not be fully controlled due to a technicality relating to constraints in the heat pump control systems.

## 3.4 Project Scope

The present study builds on the knowledge of time series analysis by developing a demand forecasting algorithm for integration into the existing control system at West Whins. Instead of simply heating the thermal store at predetermined times (figure 10), this algorithm seeks to intelligently determine when to charge the thermal store based on forecasted data. The rationale was that electricity prices tend to be high during peak demand and lower during off-peak periods. Users can thus configure the system to charge when electricity is cheaper or demand is low. In essence, the study focused on devising an algorithm that offers daily domestic hot water usage predictions to guide a smart control system.

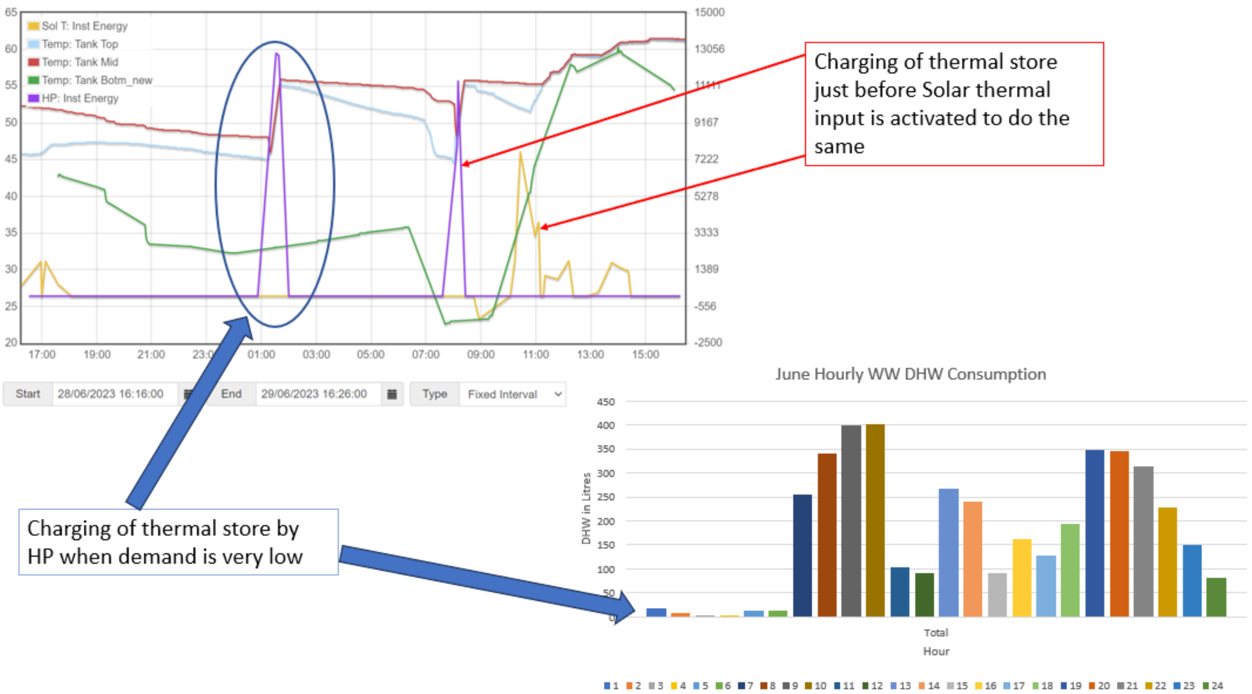


*Figure 10: Current Thermal Store charging at West Whins*

# 4.0 Literature Review

The academic research that was examined to get insight into the various methods for predicting demand for domestic hot water use in residential buildings is described in the following section.

## 4.1 Domestic Hot Water Consumption

Approximately 16-50% of the national energy consumption in many countries in residential buildings [22]. Much of this is due to domestic hot water (DHW) heating, with residential buildings in the UK using nearly 20% of domestic energy for DHW heating [23].

Forecasting hot water consumption in residential houses is a critical aspect of energy management and planning. In their study on this particular subject, Gelažanskas and Gamage, [25] explore the development of a long-term forecasting model for hot water demand in residential buildings. The authors note that hot water consumption patterns are influenced by various factors, including climate, household size, and occupants' behaviour. By analysing historical data, the researchers develop a model that incorporates these factors and captures the temporal and spatial variations in hot water consumption. The study demonstrates the importance of accurate forecasting in energy management, as it allows for better planning of energy generation and distribution resources.

## 4.2 Time Series Analysis

Time series analysis is a statistical technique used to examine data from many observations made over a long period on a single unit [26]. It is a prediction method of forecasting future values based on historical time series data [27]. Because of the difficulty in assessing the exact nature of a time series, it is usually quite challenging to generate accurate forecasts.

Demand prediction models focused on time series data can be categorised into three types based on their prediction components:

- Single models: employing a single learning algorithm;
- Ensemble models: combining multiple prediction models to determine the output data;
- Hybrid models: blending two or more machine learning techniques.

Hybrid models are particularly advantageous as they leverage the strengths of incorporated techniques resulting in enhanced forecasting accuracy and greater robustness compared to the other models [28].

Figure 11 illustrates the overall framework of time series prediction methods. Domestic hot water usage is largely dependent on time therefore this work focused on single-model time series forecasting methods for demand prediction. Subsequent sections describe the key single model time series methods utilised in this study.

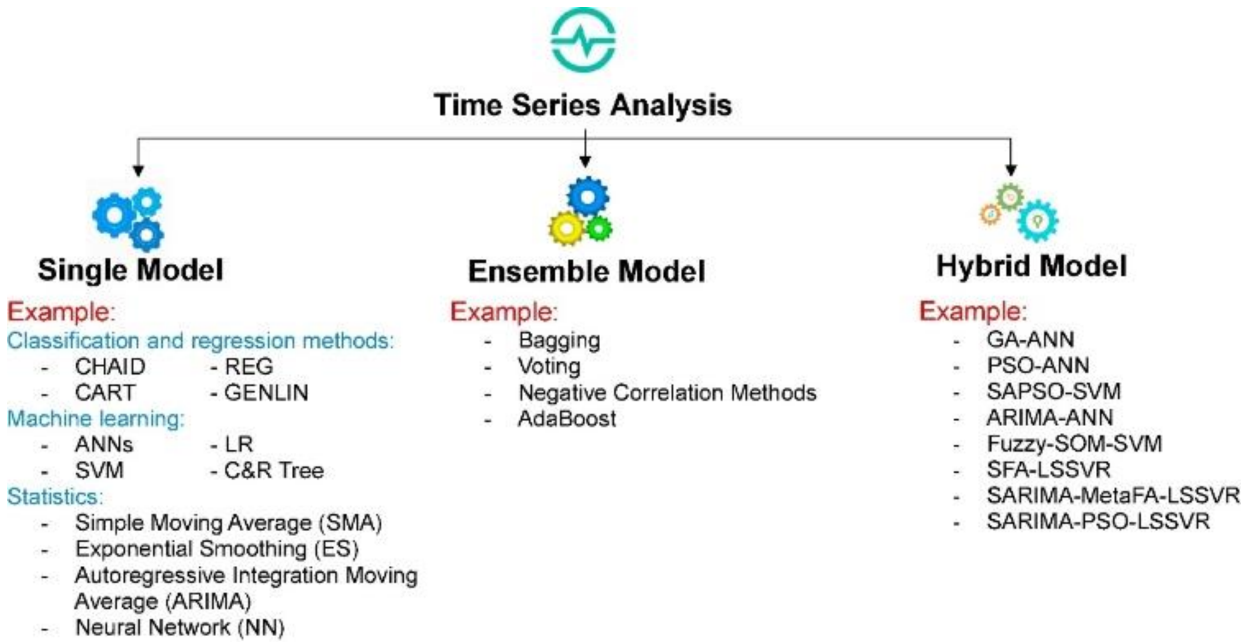


*Figure 11: Time Series Analysis*

Time series forecasting has been a fundamental area of research in the field of statistics and forecasting for many years and studies by various authors [29], [30], [31], [27], [32] provide a comprehensive overview of the advancements made in this domain in the past. The authors discuss various topics ranging from traditional univariate methods to more recent multivariate techniques. They emphasise the importance of model selection, evaluation, and validation in time series forecasting. The authors also discuss the introduction of non-linear models, such as neural networks and support vector machines (SVM), which have shown promising results in accurately predicting time series data.

In their various reviews on time series forecasting techniques for energy consumption, the researches [32], [33], [34], [35], [36], [37] all discuss three categories of forecasting methods namely statistical models, machine learning algorithms, and hybrid models. Regarding statistical models, the authors cite the use of Autoregressive Integrated Moving Average (ARIMA) and seasonal decomposition of time series (STL) as popular techniques. ARIMA models capture the dependencies of the current value on past values and the errors, while STL decomposes the series into trend, seasonal, and residual components. According to [33] ARIMA and STL have been widely used in building energy forecasting due to their simplicity and effectiveness.

In addition to statistical models, [33], [38], [39], [28] discuss the application of machine learning algorithms in energy consumption forecasting including support vector regression (SVR), artificial neural networks (ANN), and random forest (RF) as commonly used techniques. SVR is based on the principle of structural risk minimisation and has shown promising results in energy forecasting. ANN, inspired by the human brain, can capture complex nonlinear relationships in the data. RF, an ensemble learning method, combines multiple decision trees to improve prediction accuracy.

Furthermore, [40], [33], [41], [42] discuss hybrid models that integrate both statistical and machine learning techniques. These models aim to leverage the strengths of each approach to improve forecasting accuracy. The authors mention the combination of ARIMA and ANN, as well as ensemble methods such as stacking and boosting, as examples of hybrid models. According to the authors, hybrid models have shown promising results in building energy consumption forecasting by capturing both temporal dependencies and complex nonlinear relationships.

Many studies also address the issue of forecast evaluation and comparison. Authors such as [26],[43],[44],[30],[45] all present various evaluation metrics, such as mean absolute error and mean squared error, to assess the performance of different forecasting models. They also emphasise the importance of out-of-sample validation to ensure the generalisability of the models.

Walpert and Macready discuss the "No-Free-Lunch Theorem" [46] and make assertion that no forecasting method is best for every time series. Essentially, data analysts must select a forecasting method from one of the three families of forecasting techniques namely; machine learning, statistical models, and hybrid methods [36]. In this study, a forecasting system based on the statistical ARIMA technique—known to be comparatively straightforward and to perform well—was developed and assessed. The sections that follow discuss the ARIMA approach.

## 4.3 ARIMA Modelling

A time series is a sequence of observations taken sequentially in time and a normal machine-learning dataset is a collection of observations. If the current time is defined as t, then an observation at the current time is denoted by the quantity obs(t). Observations made at prior times, called lag times are denoted by the quantities obs(t-1), obs(t-2) etc where t-1 refers to the time before t and t-2 to the time before t-1. Future times are denoted by obs (t+1), obs(t+2) etc where t+1 is the time after t and t+2 the time after t+1. [31].

Autoregressive Integrated Moving Average (ARIMA) Models, developed by Box and Jenkins [17] are a popular and effective statistical technique for forecasting time series data [40]. They constitute a widely used class of linear models for the development of univariate time series prediction information [27].

ARIMA modelling is the improvement of the simpler Autoregressive Moving Average (ARMA) modelling as it adds the integral component. The key aspects of this model include the following and are summarised in the figure below:

### 4.3.1 Autoregression (AR)

In an autoregression model, the variable of interest is predicted using a linear combination of the variable's prior values. It is therefore a regression of the variable against itself. Using standard notation an autoregressive model of order p can be written as:

$yt = C + \alpha 1 yt - 1 + \alpha 2 yt - 2 + \cdots + \alpha p yt - p$

where C is a constant; $yt-1,yt-2,yt-p$ are the lags (past values); and $\alpha1,\alpha2,\alpha p$ are lag coefficients which are estimated by the model.

An autoregression model assumes that the observations at current and previous time steps are useful to predict the value at the next time step. This relationship between variables is called correlation.

#### 4.3.2 The I (Integrated) term

This term refers to the use of a differencing operator of raw observations (e.g., subtracting an observation from an observation at the previous time step) to make the time series stationary.

#### 4.3.3 Moving Average

The moving average part of the ARIMA model describes the dependency between the current observation and a residual error from a moving average model applied to lagged observations. Using standard notation, the moving average yt can be described by the function:

$yt=\epsilon t+\beta1\epsilon t-1+\beta2\epsilon t-2+\cdots+\beta q\epsilon t-q$

Where:

$\epsilon t,\epsilon t-1,\epsilon t-q$ are white noise terms for the respective lags, i.e., $yt-1,yt-2,yt-q$; and $\beta1,\beta2,\beta q$ are the parameters of the model [31].

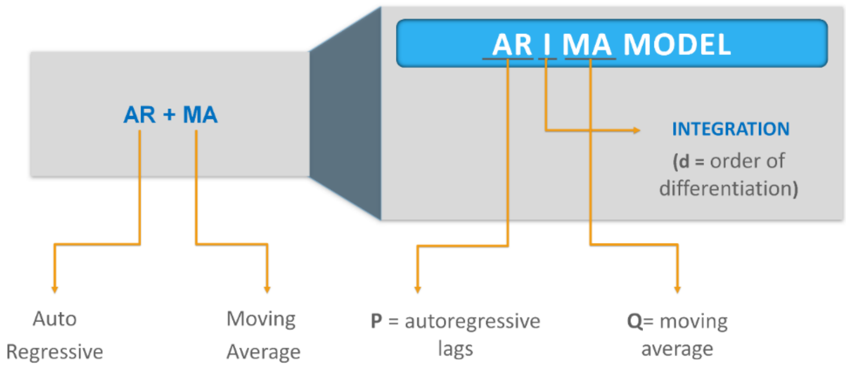


*Figure 12: ARIMA Model Summary*

These parts are explicitly specified in the model as a parameter. A standard notation that takes the form ARIMA(p,d,q) is used for this, where the parameters p, d and q are substituted with finite integer values to indicate the specific ARIMA model being used. The parameters of the ARIMA model are defined as follows:

1. ‘p’ is the number of autoregressive terms or number of lag observations in the model, also known as the lag order;
2. ‘d’ is the number of differences or nonseasonal differences needed for stationarity; and
3. ‘q’ is the number of moving averages or lagged forecast errors in the prediction equation[47], [38].

The constants p and q are the model orders while the term d represents the degree of ordinary differencing, applied to make the series stationary [40]. The appropriate orders of the ARIMA(p, d, q) model are usually determined through the Box-Jenkins model methodology [27].

## 4.4 Box-Jenkins Method

The Box-Jenkins method was proposed by George Box and Gwilym Jenkins in their seminal 1970 textbook Time Series Analysis: Forecasting and Control [48]. Box and Jenkins based their approach on the assumption that the process that generates a time series can be approximated using an ARIMA Model if it is stationary. Model identification, parameter estimation, and diagnostic verification are the three iterative processes of the Box-Jenkins approach [17].

The identification stage is meant to first asses if the time series data is stationary and if not how many differences (d) would be required to make it stationary. Box and Jenkins further suggested producing two diagnostic plots using the sample's data to choose the p and q parameters of the ARIMA model. These are:

1. Autocorrelation Function (ACF) – a plot to summarise the correlation of an observation with lag values. The x-axis shows the lag and the y-axis shows the correlation between the -1 and 1 for negative and positive correlation.

2. Partial Autocorrelation Function (PACF) – a plot of the summary of the correlations for observation with lag values that are not accounted for prior lagged observations.

Parameter estimation involves using numerical methods to minimise a loss or error term, while diagnostic checking looks for evidence that the model is not a good fit for the data by using a review of residual errors. The errors from an ideal model would resemble white noise, that is a Gaussian distribution with a mean of zero and a symmetrical variance. Plots such as a density plot, histograms and Q-Q plots are used to compare the distribution of errors. An asymmetrical distribution or a non-zero mean may suggest a bias in the forecasts produced.

Further diagnostic checks can be carried out by creating more ACF and PACF plots of the residual error time series and the presence of a serial correlation would suggest further opportunity for model analysis. This three-step model creation method is often repeated multiple times until a good model is finally chosen. The final model chosen can then be applied to make predictions.

The following table summarises the broad categories of demand forecasting techniques and their uses.

*Table 1: Summary of Literature Review*

| Broad Category | Specific Method | Findings/Use | Source |
|---|---|---|---|
| Machine Learning Models | Support Vector Machine (SVM) models with time series clustering | Forecasted water demand in a distribution network with a 24-hour delay | [49, 50] |
| | Long-Short term memory network (LSTM) | Compared to ARIMA, SVM, and random forests models, LSTM showed better results | [51] |
| | ARIMA models, Optimized Theta, ensemble methods, and neural networks | Forecasted water consumption with a horizon of several months. Concluded that ARIMA models performed best based on various metrics. | [52] |
| | Neural networks to predict domestic hot water (DHW) consumption in residential buildings | Observed better prediction performance as the size of the systems increased | [53] |
| | Autoregressive neural network | Demonstrated the importance of external variables for improving predictions | [54] |
| | Seasonal ARIMA (SARIMA) | Analysed the effect of seasonality on regression performances for DHW load prediction | [25] |
| | Artificial neural networks (ANN) | Benefits include continuous adaptability to changing water demand patterns and applicability to different demand signals | [55, 56] [57] |
| | Random forests (ensemble) | Success was attributed to a good choice of features in models | [58] |

| | | | |
|---|---|---|---|
| Statistical and Periodic based Models | Multilinear and nonlinear regression | The method is based on data-driven techniques | [59] |
| | Regression techniques | A pattern-based approach to avoid problems encountered with next-day transitions | [60] |
| | Holt-Winter, ARIMA, and GARCH | Statistical forecasting for pattern recognition | [61] |
| | Moving window | Demand fluctuations are determined by variables calibrated on a shifting window of observed data | [62]<br>[63] [64] |
| Stochastic Models | Bayesian model | Analysed uncertainty quantification and reduction | [65] |
| | Model conditional processor (MCP) | Combined different forecasting models and estimated predictive uncertainty | [66] [42] |
| | Markov chain-based model | Estimated the probability that future demands will fall within given ranges | (Gagliardi et al. 2017) |

## 4.5 Metrics for Evaluation

Time series prediction performance metrics give an overview of the knowledge and expertise of the forecast model that produced the predictions. There is a wide variety of performance metrics available for this. Metrics can be used to assess the model's performance and they have been shown to depict its accuracy very well. They are computed for each model and compared with each other to identify the most accurate one. The following table summarises these metrics:

*Table 2: Performance Metrics*

| Error Type | Description | Formula |
|---|---|---|
| Forecast Error | This is the expected value minus the predicted value called the residual error of the prediction. A forecast error of zero indicates no error or perfect skill for that forecast. | forecast error = expected value − predicted value |
| Mean Forecast Error | This is the average of the forecast error values and an ideal mean forecast error would be zero. A mean forecast error value other than zero suggests a tendency of the model to over-forecast (negative error) or under-forecast (positive error). A forecast bias of zero, or a very small number near zero, shows an unbiased model. | mean forecast error = mean(forecast error) |
| Mean Absolute Error (MAE) | MAE is the average of the forecast error values, where all of the forecast values are forced to be positive. A mean absolute error of zero indicates no error. | mean absolute error = mean(abs(forecast error)) |
| Mean Squared Error (MSE) | MSE is calculated as the average of the squared forecast error values. A mean squared error of zero indicates perfect skill or no error. | mean squared error = mean(forecast error2) |
| Root Mean Squared Error (RMSE) | The square root of the mean squared error score is called the root mean squared error, or RMSE. As with the mean squared error, an RMSE of zero indicates no error. | rmse = p mean squared error |

## 4.6 Conclusion and Gap Statement

In the 2023 article "Futures for Findhorn" by Copeland et. al.[67], challenges faced by the Findhorn community in achieving a net-zero carbon footprint are highlighted. At the time of writing, despite their sustainability efforts, they struggle with issues like significant energy imports. The community's energy consumption, particularly in heating and transportation, heavily influences their emissions. The authors emphasise a comprehensive approach that includes technology, behaviour changes, and community engagement, with a focus on energy efficiency, renewables, and green transport.

The current study employed the ARIMA method for time series analysis, given its proven efficacy in forecasting using past data. Data collected hourly and half-hourly from the Findhorn ecovillage validated the ARIMA model's suitability for predicting DHW energy demand.

# 5.0 Methodology

The approach, technical analysis, and modelling used to accomplish the aims and objectives are described in this section. Subsequent subsections provide an explanation of model development techniques used.

The table below lists the hardware and software specifications of the system used in the development of the models for the current work.

*Table 3: Hardware and software specifications*

| | **Details** |
|---|---|
| **Processor** | 11th Gen Intel(R) Core(TM) i7-1165G7 @ 2.80GHz 2.80 GHz |
| **Storage** | 1TB SSD; Installed RAM - 16.0 GB (15.7 GB usable) |
| **Display** | Intel® Iris® Xe Graphics, MS Excel, Lucidchart |
| **Software** | Python 3.11.4, VS Code (Jupyter Notebook 6.5.2), Anaconda (Spyder 5.4.1) |
| **Libraries** | Pandas, numpy, matplotlib, statsmodels, sklearn.metrics, scipy.stats |

## 5.1 Project Approach

The following flow chart provides a visual representation of the integrated process undertaken during the project.

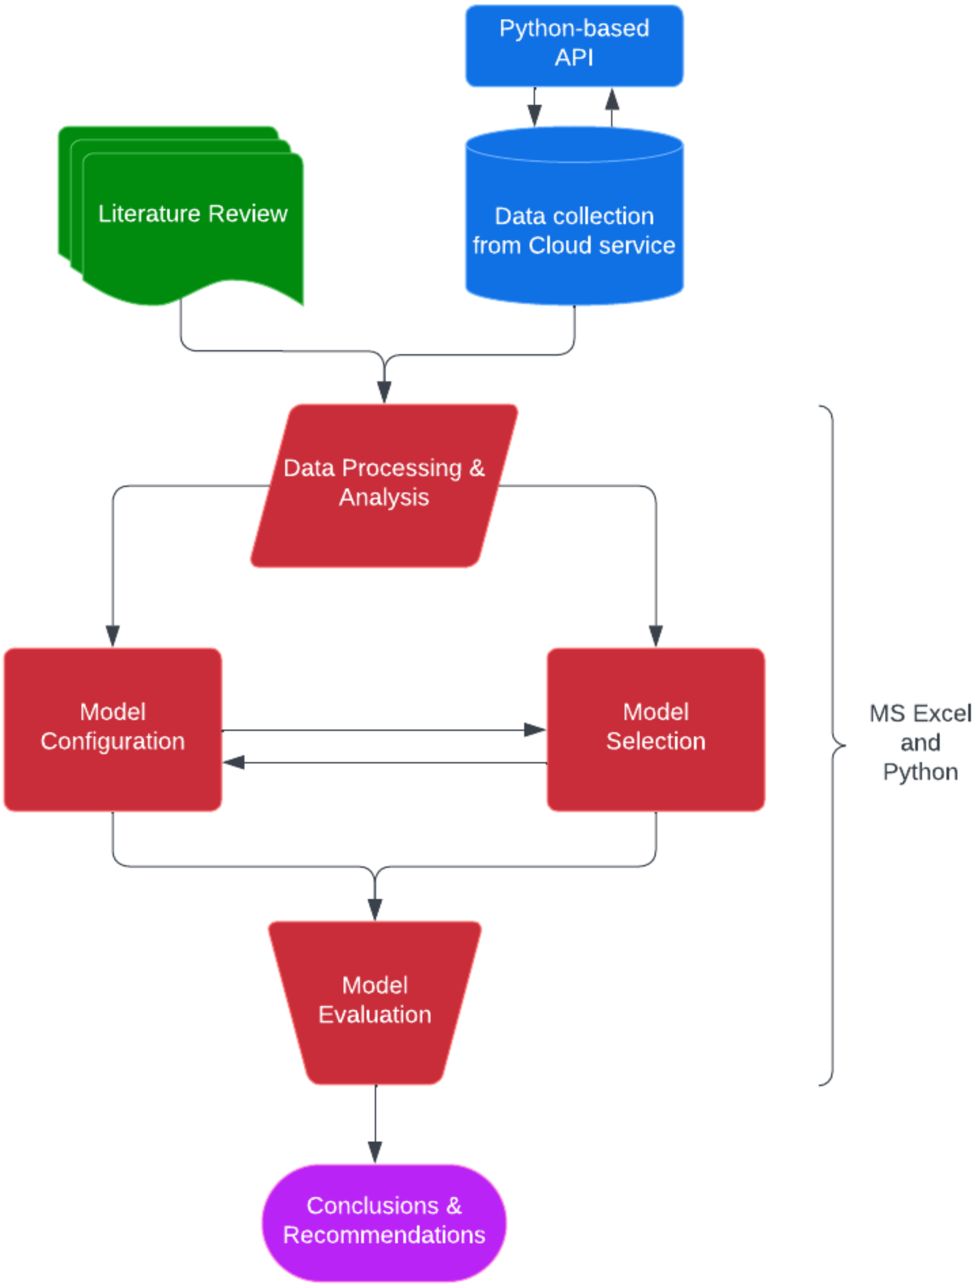


*Figure 13: Project Approach*

## 5.2 Time Series Model Development

To develop a time series forecasting model for predicting DHW consumption, various procedures were undertaken, which are detailed in this section. Drawing from methodologies outlined by authors [68] and [69], a procedure was formulated to predict DHW energy usage for both summer and winter periods.

## 5.3Data Collection and Visualisation

The time series data used in the model development was collected from the Emoncms cloud service using a Python-based API. The data was cleaned to remove any outliers, inconsistencies or missing values. It was further rearranged such that the averages of the hourly consumptions for each time period (winter and summer) were obtained and saved in separate excel datasets in the model development file location.

In time series analysis and forecasting, visualisation plays an important role. Plotting the configured datasets helped to uncover time-based characteristics like trends, cycles, and seasonality that can influence the model selection.

## 5.4 Baseline Model

A baseline in forecast performance provides a point of comparison. It was important to gauge the performance and reliability of any elaborate models developed later by establishing a clear baseline for comparison. The chosen model for comparison was a Simple Moving Average Model (SMA) which is a method used to identify trends by smoothing out large fluctuations. This is done by calculating the average of the data points within a specific window of periods. The level of accuracy of the baseline model was evaluated using metrices, the RMSE and the MAE.

## 5.5 Data Stationarity

A stationary time series is one where the values are not a function of time. Time series are stationary if they do not have trend or seasonal effects and summary statistics calculated on the time series are consistent over time, like the mean or the variance of the observations. Experts suggest that a stationary time series is easier to model and statistical modelling methods assume or require the time series to be stationary to be effective.

The stationarity of the learning datasets was confirmed using a statistical significance test, specifically the Augmented Dickey-Fuller (ADF) test, which was performed by running a script with the adfuller() function in the Statsmodels Python library. It uses an autoregressive model and optimizes an information criterion across multiple different lag values. The null hypothesis of the test is that the time series can be represented by a unit root and that it is not stationary (has some time-dependent structure). The alternate hypothesis (rejecting the null hypothesis) is that the time series is stationary.

- *Null Hypothesis (H0)*: Fail to reject, it suggests the time series has a unit root, meaning it is non-stationary. It has some time-dependent structure.

- *Alternate Hypothesis (H1)*: The null hypothesis is rejected; it suggests the time series does not have a unit root, meaning it is stationary. It does not have a time-dependent structure.

The result is interpreted using the p-value from the test. A p-value below a threshold (such as 5% or 1%) suggests we reject the null hypothesis (stationary), otherwise, a p-value above the threshold suggests we fail to reject the null hypothesis (non-stationary).

- *p-value > 0.05*: Fail to reject the null hypothesis (H0), the data has a unit root and is non-stationary.

- *p-value ≤ 0.05*: Reject the null hypothesis (H0), the data does not have a unit root and is stationary [31].

Where the time series was found to be non-stationary, data transformation would be required to make the time series stationary.

## 5.6 ARIMA Model Configuration

Autocorrelation Function (ACF) and Partial Autocorrelation Function (PACF) plots were generated in Python to determine initial model parameters. To better determine the optimal ARIMA configuration (the p, d, q order), various combinations were systematically explored. During this process, two key metrics were used, the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC). Both metrics evaluate model goodness of fit, with a penalty for models that might be too intricate.

Essentially, lower AIC or BIC values denote models that strike an ideal balance between accuracy and simplicity. Using Python, these metrics were generated after model fitting with a goal to select the model showcasing the lowest AIC and BIC scores.

## 5.7 Model Evaluation Using Residual Errors

To further evaluate the accuracy of the model, a review of the residual forecast errors was carried out. Residual errors, in essence, are the differences between predicted and actual values. Ideally, the distribution of residual errors should be random and normally distributed around zero. If there's a pattern in the residuals, it suggests that the model may not be capturing some information.

The following steps were implemented in Python to evaluate the ARIMA model using residual errors:

1. Compute Residuals: Subtract the predicted values from the actual values.
2. Plot Residuals: This helps in visually checking if the residuals have any patterns or seasonality.
3. Check for Mean Residual: Ideally, the mean of the residuals should be zero or very close to zero.
4. Plot Residual Distribution: Use histograms or kernel density plots to verify if residuals are normally distributed.
5. Residual Autocorrelation: Use the ACF (Autocorrelation function) plot to check if residuals have any correlation with lagged versions of itself.

The performances of the developed models were evaluated and conclusions drawn based on the analysis of the associated errors.

### 5.8 High-Level Python Script Process

Figure 14 provides an overview of the algorithm used to forecast energy consumption, illustrating how the functions (.py) within the forecasting script produce the desired outcomes.

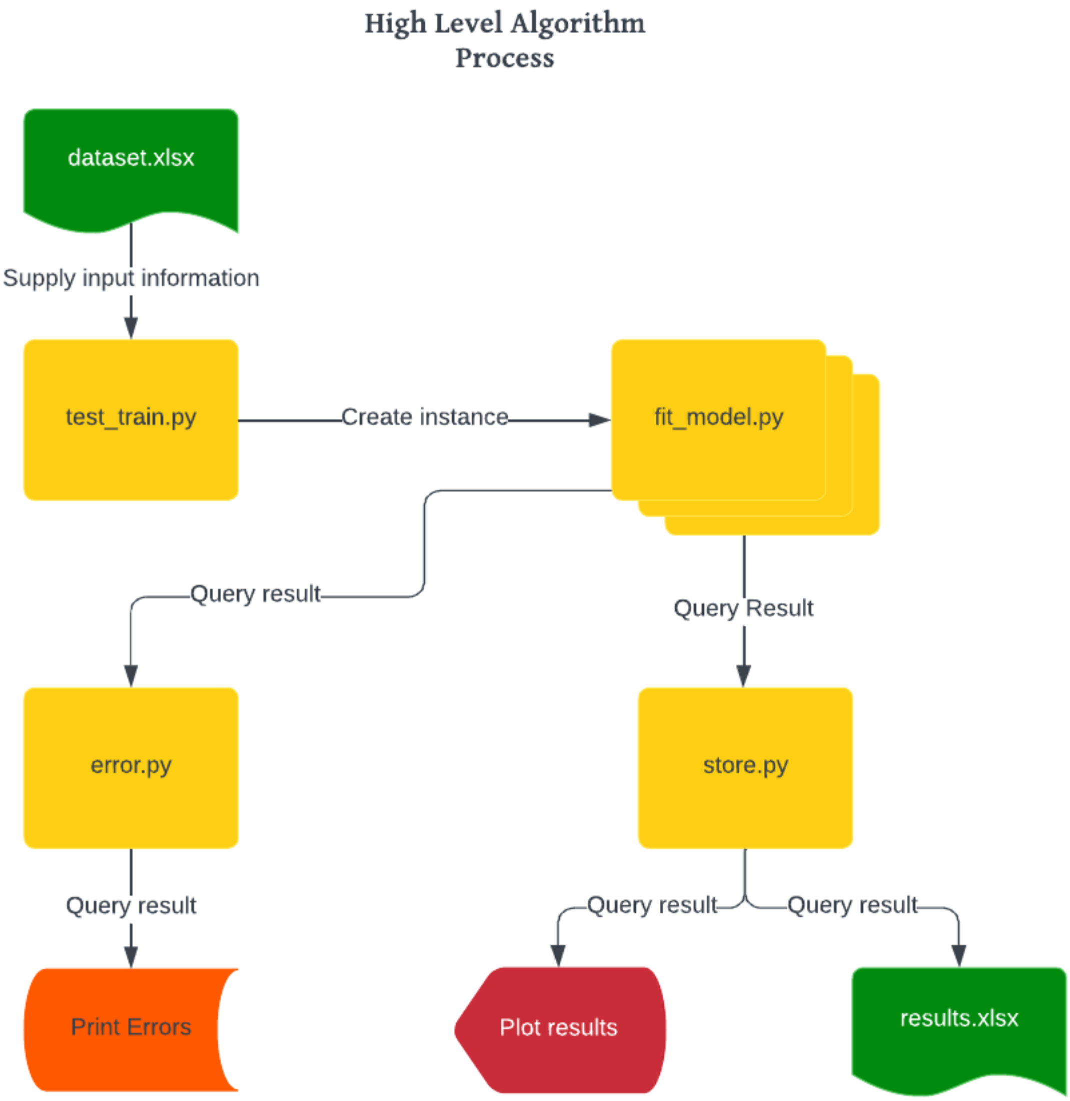


*Figure 14: High Level Algorithm Process*

## 6.0 Technical Results Analysis

The following section presents the technical results gathered during the study. An analysis of the findings is also provided.

### 6.1 Data Analysis using summary statistics and plots

Summary statistics and plots analysed the data's structure and characteristics, focusing on daily and hourly volumetric hot water consumption and its energy delivery.

Hot water consumption was measured to a 0.1 litre resolution from the six dwellings in the West Whins residential building. The plots and tables summarise the data collected for one summer and one winter season (91 days each).

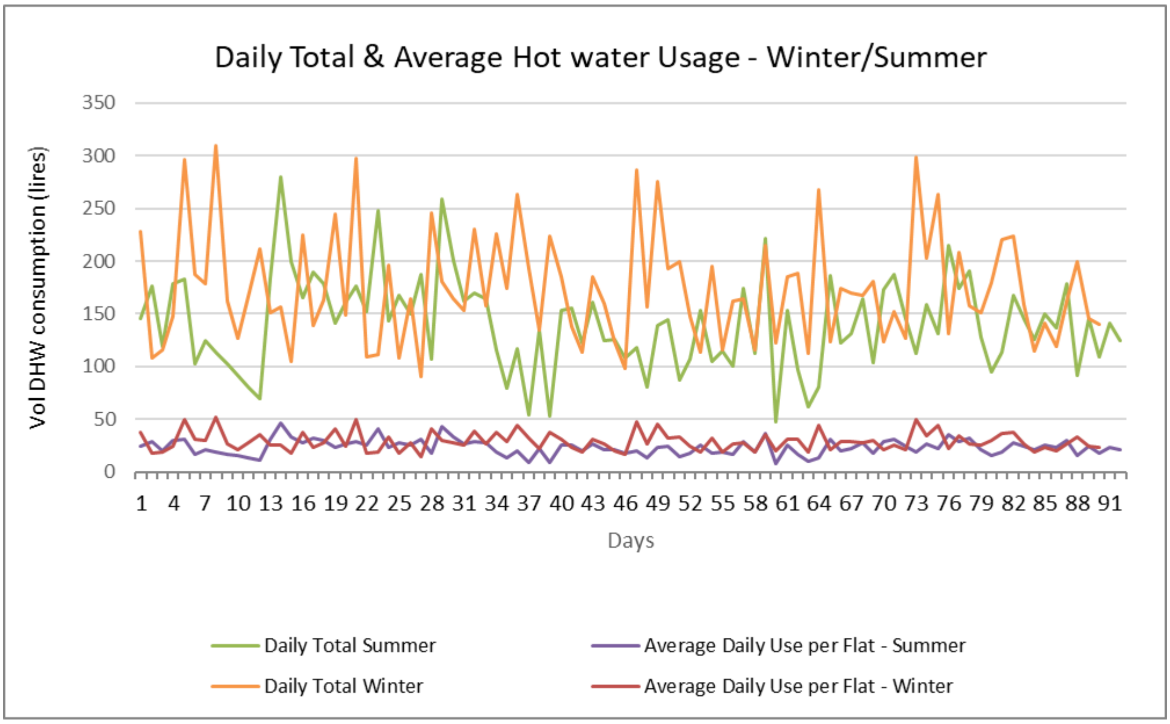


*Figure 15: Daily Total against Daily Average Usage for Winter and Summer*

*Table 4: Daily Water Consumption – Litres/day*

| | ***Winter*** | | ***Summer*** | |
|---|---|---|---|---|
| | Daily Total Building Consumption | Daily average consumption per flat | Daily Total Building Consumption | Daily average consumption per flat |
| ***Mean*** | 174 | 29 | 140 | 23 |
| ***Max*** | 309 | 52 | 280 | 47 |
| ***Min*** | 90 | 15 | 48 | 8 |

The average hourly hot water consumptions per flat and for the whole building are shown below (Figures 16 - 17) as well as the cumulative hourly consumption per season (Figure 18).

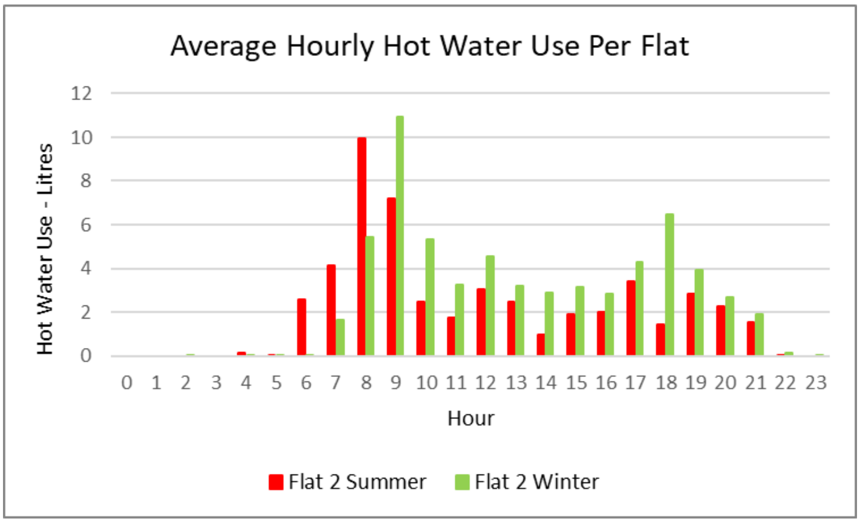


*Figure 16: Average Hourly Water Usage for one dwelling*

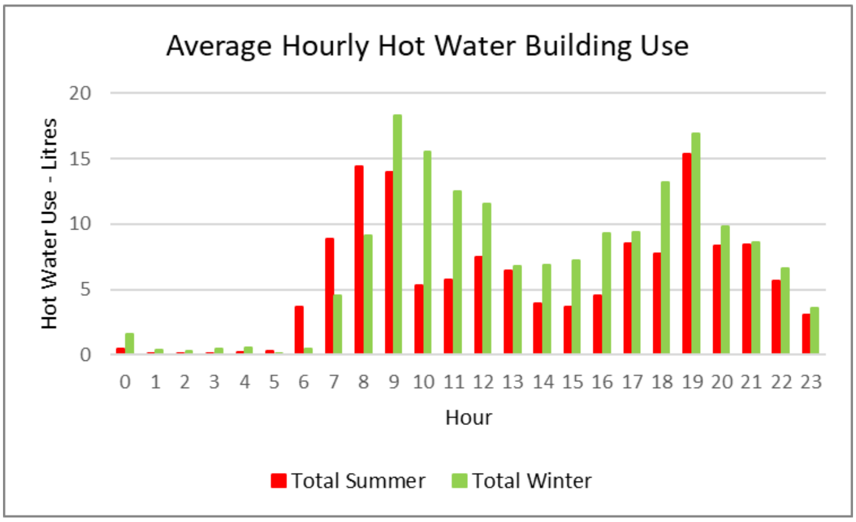


*Figure 17: Average Hourly Water Usage for the whole building*

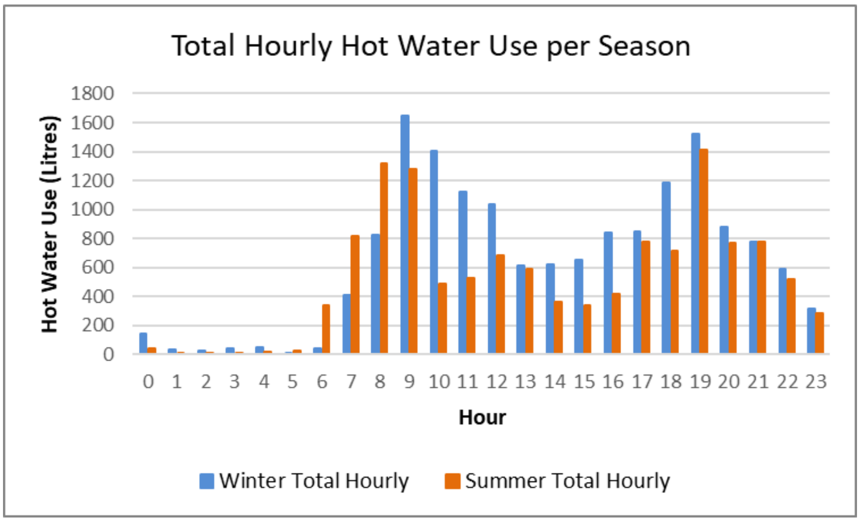


*Figure 18: Cumulative Hot water use per season*

The plots revealed a consistent pattern of Domestic Hot Water (DHW) usage across both daily and hourly consumption for both the summer and winter seasons. This consistency suggested that the factors driving DHW consumption in the building were relatively stable over time.

Moreover, the absence of significant outliers in the plots underscored the reliability of the data collection process and the normalcy of the DHW usage patterns observed. Such consistent patterns and the lack of extreme outliers support the application of time series modelling techniques for future DHW consumption prediction.

### 6.1.1 Energy Consumption

To determine the energy content of the hot water used, the volumetric water consumption and the temperature rise required to heat it was used in the general energy equation:

$$Q = mc\Delta T$$

Where Q is the heat energy (J or kWh)
m is the mass of the water (kg)
c is the specific heat capacity of water (~4200 J/kg°C)
ΔT is the temperature change (°C)

The energy content was calculated based on the volumetric hot water consumption values (assuming 1L of water equals 1kg) and the nominal temperature difference of 43°C required to heat the water, as stated previously, and a conversion of energy in joules to kwh. Figure 19 depicts the daily energy consumption of the building.

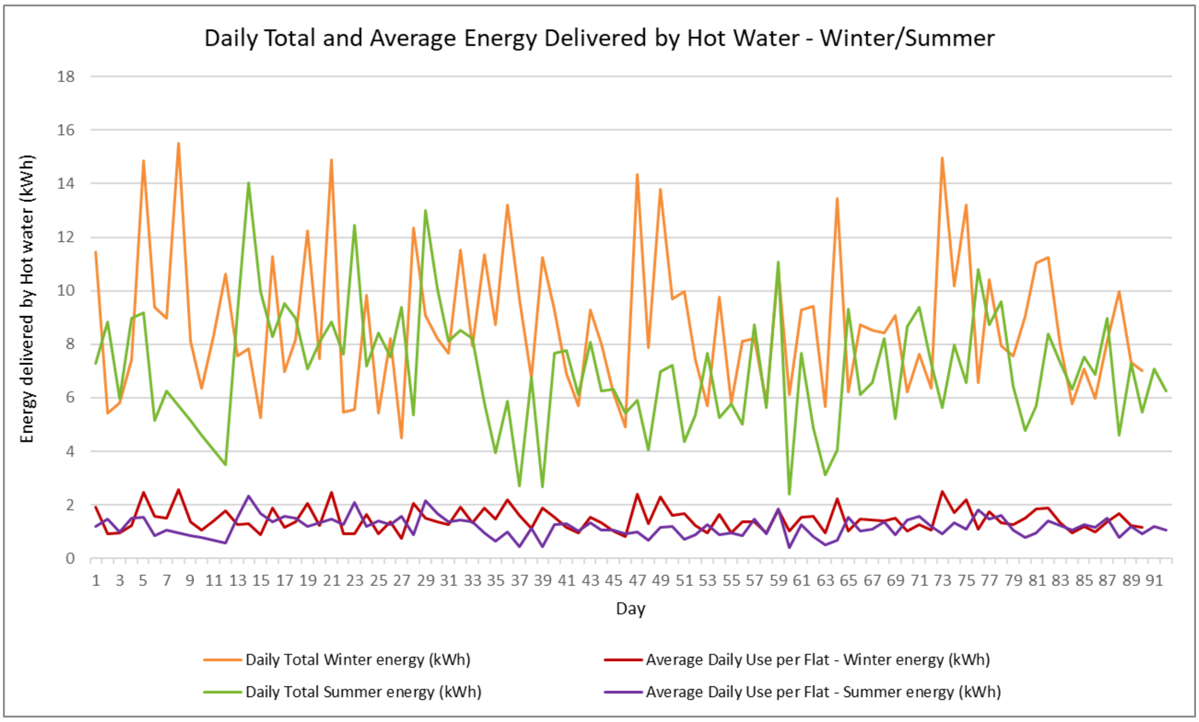


*Figure 19: Daily distribution of energy delivered across the whole sample*

A summary of the daily energy consumption statistics is given in the table below:

*Table 5: Daily Energy Distributed in Hot Water Use - kWh*

| | ***Winter*** | | ***Summer*** | |
|---|---|---|---|---|
| | Building Energy Consumption | Daily Energy Consumption per flat | Building Energy Consumption | Daily Energy Consumption per flat |
| ***Mean*** | 8.72 | 1.45 | 7.05 | 1.17 |
| ***Max*** | 15.5 | 2.58 | 14.05 | 2.34 |
| ***Min*** | 4.52 | 0.75 | 2.41 | 0.40 |

Hourly energy distribution to hot water was plotted for each of the 6 flats in both winter and summer periods (Figure 20 - 22).

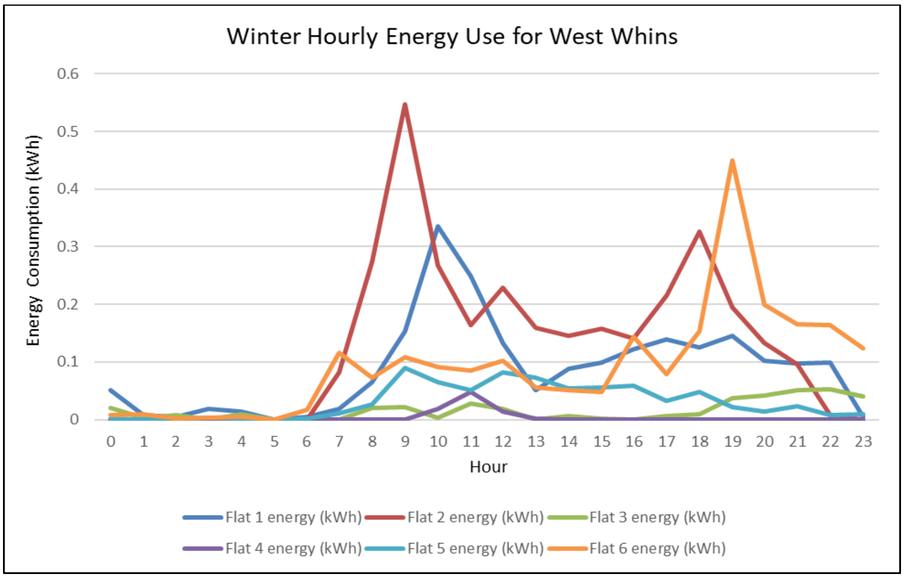


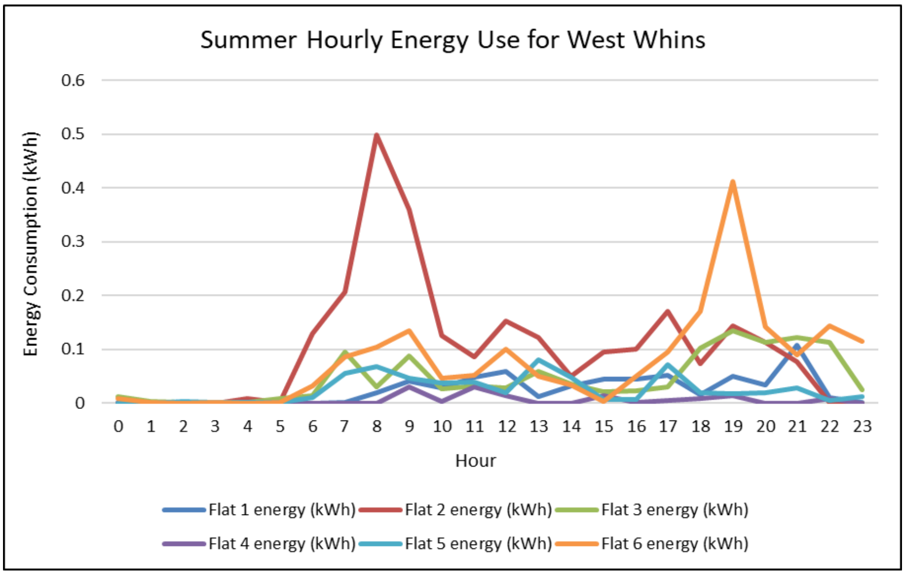


*Figure 20: Comparison of Hourly Energy Usage Per Flat at West Whins*

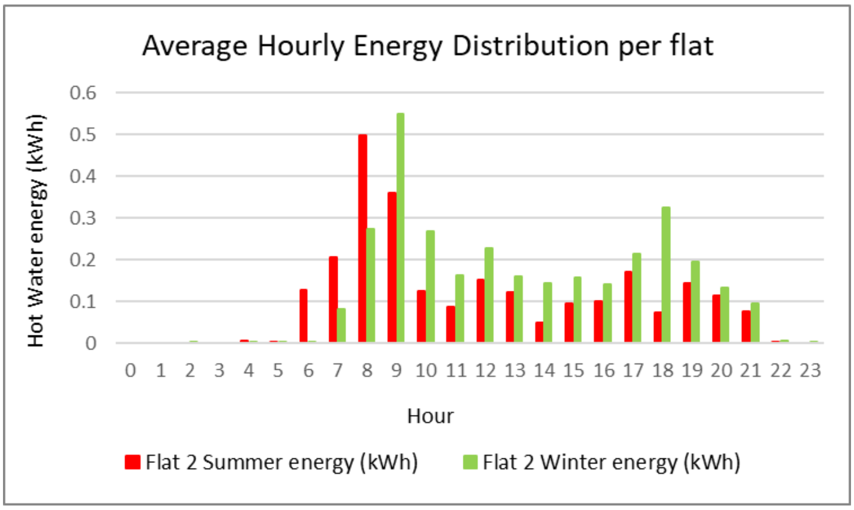


*Figure 21: Hourly energy distribution delivered to hot water in one dwelling*

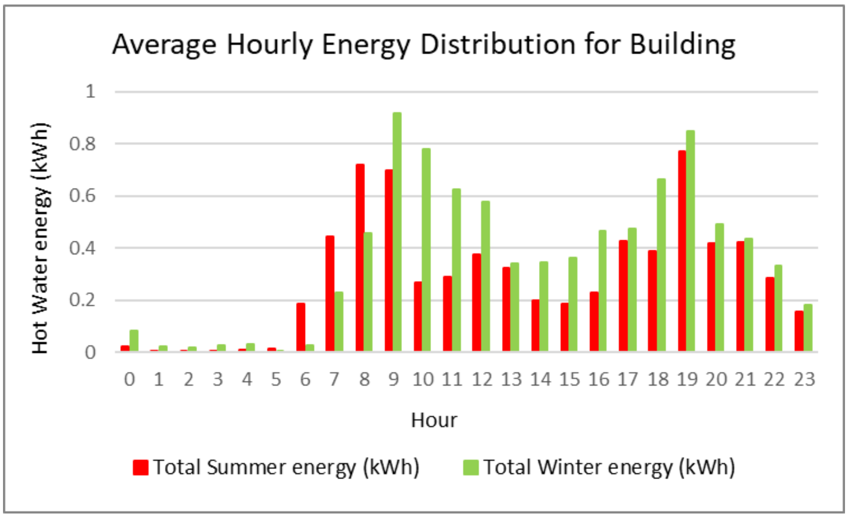


*Figure 22: Hourly energy distribution delivered to hot water for the whole building*

## 6.2 Creation of baseline models

A Simple Moving Average (SMA) was created as a baseline model. This was done using a python script and the RMSE and the MAE evaluation metrices were computed.

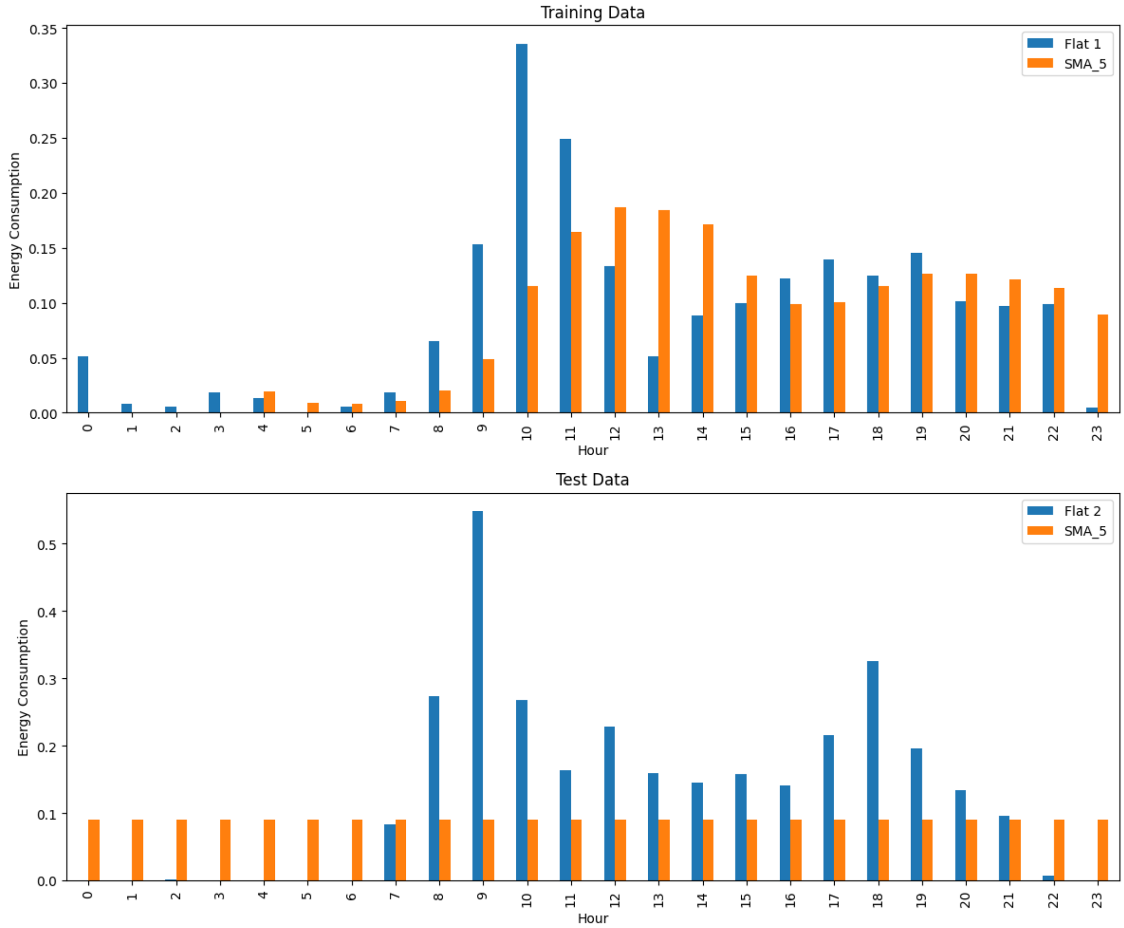


*Figure 23: Simple Moving Average Baseline Model Outputs – Winter*

Using a Simple Moving Average (SMA) model with a window of 5 hours, the energy usage for Flat 2 was predicted based on data from Flat 1.

*Table 6: Simple Moving Average Evaluation Metrices*

| ***Winter Stats*** | *Training Data* | *Test Data* |
|---|---|---|
| *RMSE* | 0.0734 | 0.1400 |
| *MAE* | 0.0507 | 0.1083 |

The model had an acceptable level of predictive power with a Root Mean Squared Error (RMSE) of 0.0734 and Mean Absolute Error (MAE) of 0.0507 for the training data, and an RMSE of 0.1400 and an MAE of 0.1083 for the test data. Similar results were observed using the summer datasets as summarised below.

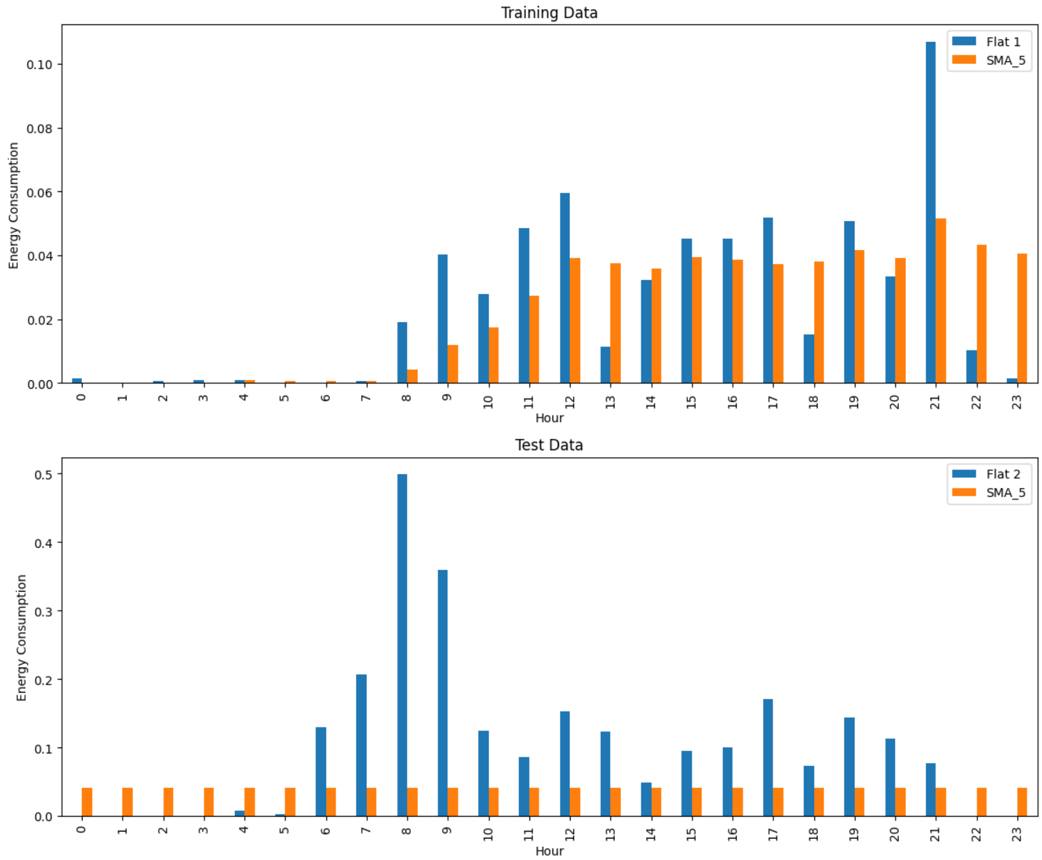


*Figure 24: Simple Moving Average Baseline Model Outputs – Summer*

*Table 7: Simple Moving Average Evaluation Metrices*

| ***Summer Stats*** | *Training Data* | *Test Data* |
|---|---|---|
| *RMSE* | 0.0215 | 0.1338 |
| *MAE* | 0.0159 | 0.0903 |

The simple baseline model predicted a constant energy usage for all hours in Flat 2 for both datasets due to the nature of the SMA model, which took the last 5 hours of the training data to forecast future usage (window size of the Simple Moving Average (SMA) model).

Rather than forecasting a uniform energy usage for all hours in Flat 2 using the last 5-hour average from the training data in Flat 1, the model was enhanced to recalculate the SMA for each test hour based on the preceding 5 hours of combined training and test data. This change allowed for hour-by-hour variation in predictions, offering a closer representation of the latest data trends. The results for the winter and summer datasets are shown below.

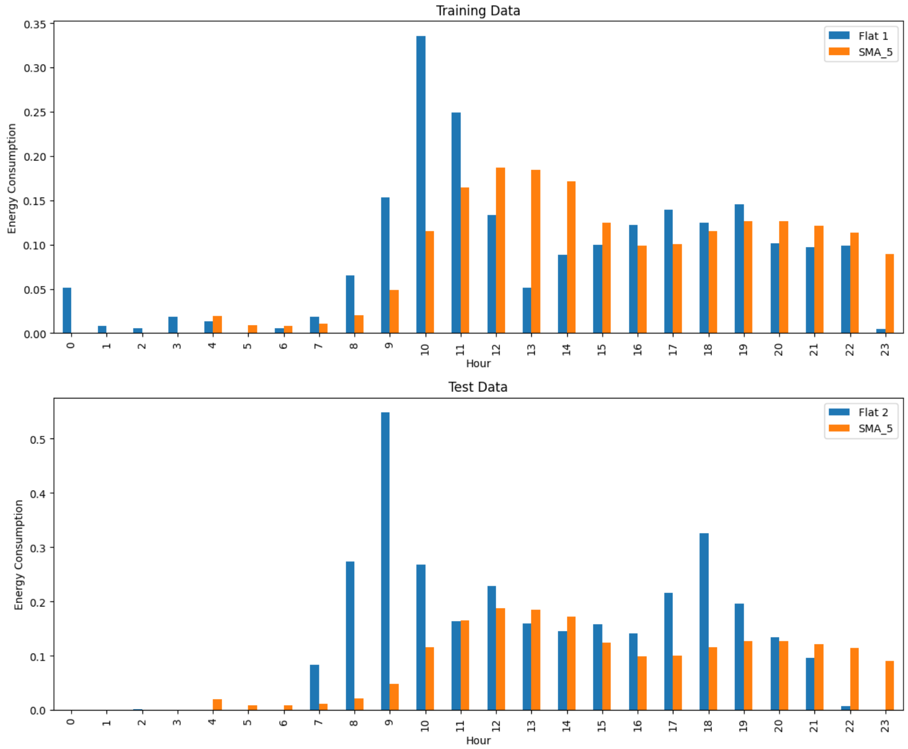


Figure 25: Enhanced SMA Baseline Model Outputs – Winter

Table 8: Enhanced SMA Evaluation Metrices

| *Winter Stats* | *Training Data* | *Test Data* |
|---|---|---|
| *RMSE* | 0.0734 | 0.1466 |
| *MAE* | 0.0507 | 0.0901 |

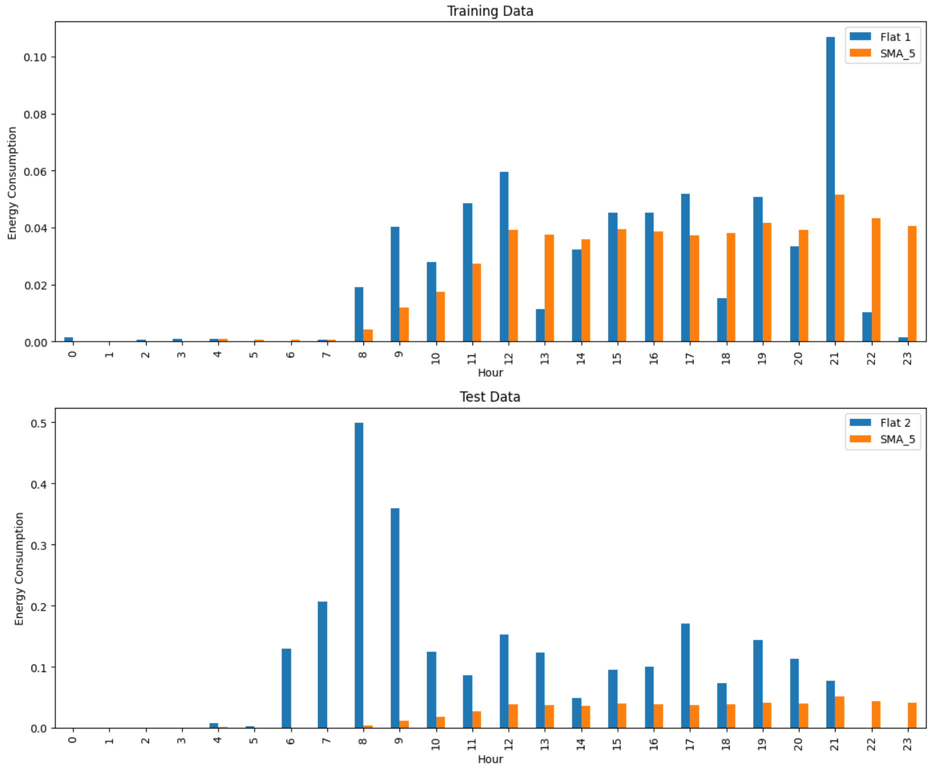


Figure 26: Enhanced SMA Baseline Model Outputs –Summer

*Table 9: Enhanced SMA Evaluation Metrices*

| ***Summer Stats*** | *Training Data* | *Test Data* |
|---|---|---|
| *RMSE* | 0.0215 | 0.159 |
| *MAE* | 0.0159 | 0.107 |

This simple model served as a baseline for future, more complex models and guided subsequent steps of model selection and tuning to improve predictive power for the buildings hourly energy consumption.

The same was done with the daily consumption predictions for both summer and winter datasets. Due to the nature of the daily consumption data, the datasets were split using the first two months of the seasons as training sets (December and January for winter and June and July for summer) and the last month as the test sets (February for winter and August for summer). The results are shown as follows:

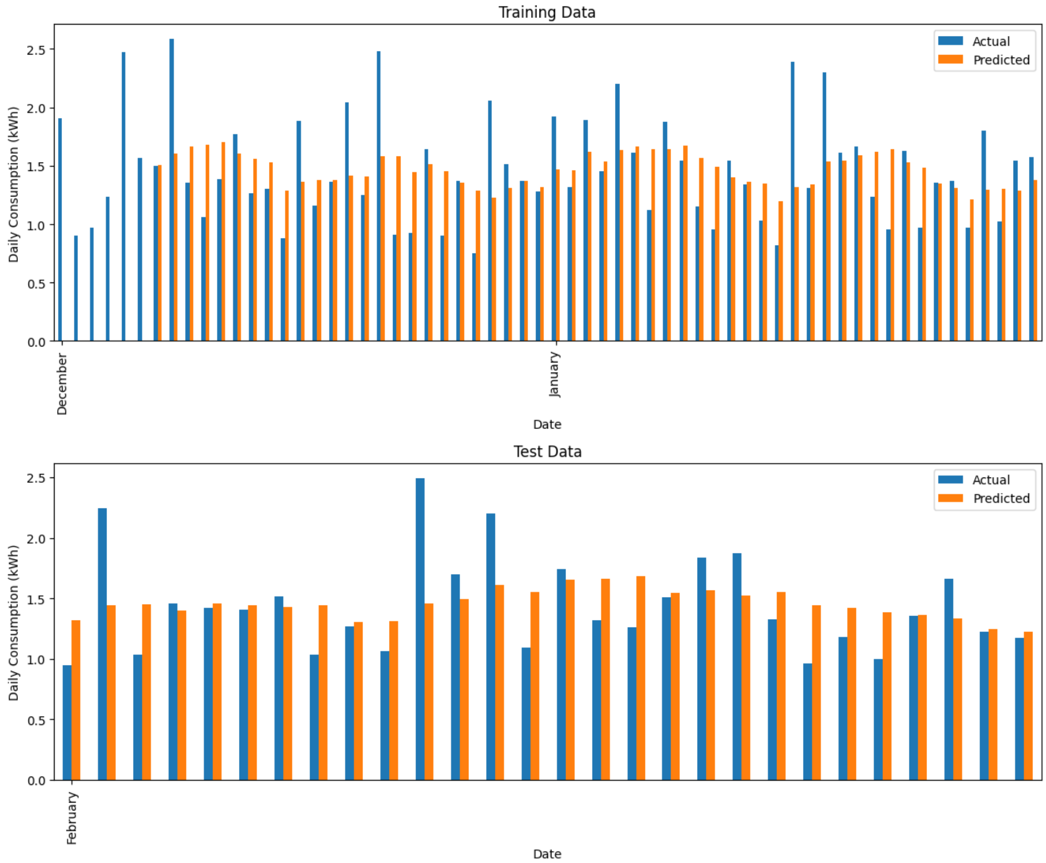


*Figure 27: Enhanced SMA Baseline Model Outputs – Winter*

*Table 10: Enhanced SMA Evaluation Metrices - Winter*

| ***Winter Stats*** | *Training Data* | *Test Data* |
|---|---|---|
| *RMSE* | 0.4295 | 0.3759 |
| *MAE* | 0.3355 | 0.2864 |

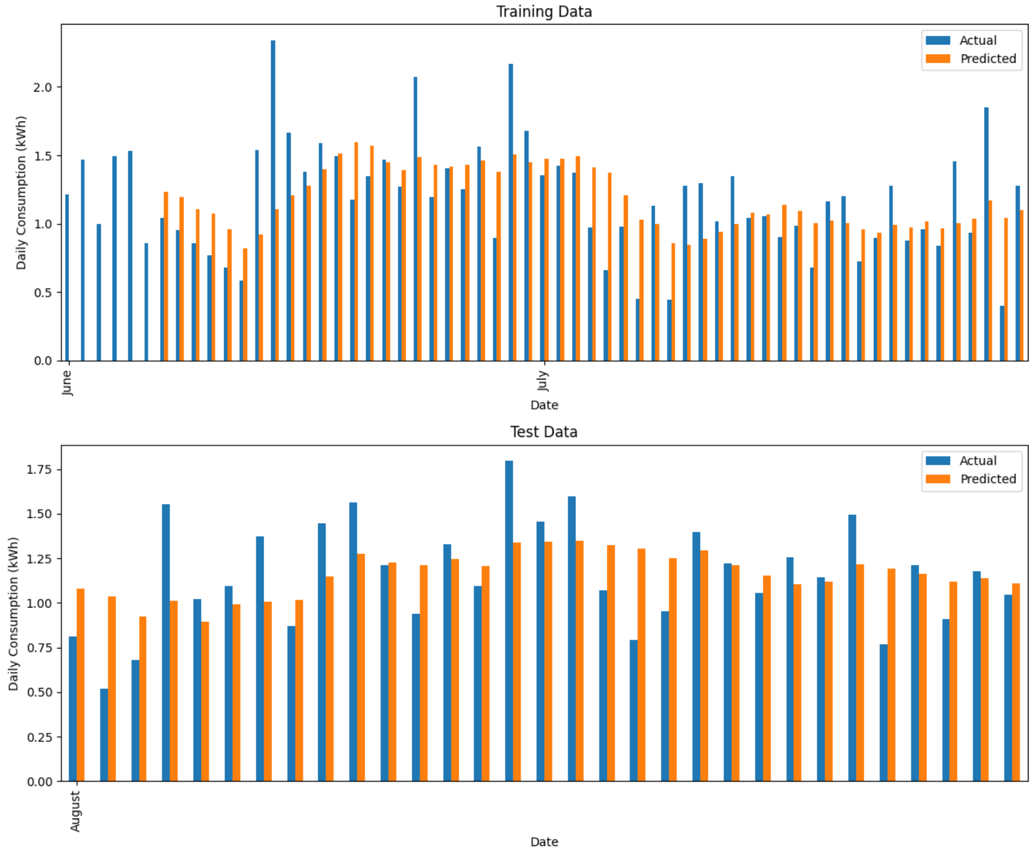


*Figure 28: Enhanced SMA Baseline Model Outputs – Summer*

*Table 11: Enhanced SMA Evaluation Metrices - Summer*

| ***Summer Stats*** | *Training Data* | *Test Data* |
|---|---|---|
| *RMSE* | 0.3634 | 0.2662 |
| *MAE* | 0.2801 | 0.2169 |

Similar to the hourly prediction data, the model using the daily data recalculated the SMA for each test day based on the preceding 5 days of combined training and test data. This allowed for daily variation in predictions, offering a closer representation of the latest data trends

## 6.3 Data Stationarity

The Augmented Dickey-Fuller (ADF) test was conducted to check data stationarity on the hourly datasets. The test returned a p-value of 0.08 for the winter dataset, which was above the 0.05 threshold, indicating potential non-stationarity. The ADF statistic of -2.66, however, was less than the 10% critical value (-2.64), suggesting possible stationarity at this level. Hence, the results presented a mixed picture regarding the data's stationarity.

*Table 12: Results of Dickey-Fuller Test:*

| | |
|---|---|
| ADF Statistic | -2.6649 |
| p-value | 0.0803 |
| Critical Value (1%) | -3.7697 |
| Critical Value (5%) | -3.0054 |
| Critical Value (10%) | -2.6425 |

In order to apply ARIMA modelling, which requires stationary data, a first-order-differencing method (d = 1) was performed on the energy consumption data from Flat 1. The augmented Dickey-Fuller test was then carried out on the transformed data to assess its stationarity. The results of the test are shown in figure 28.

a b

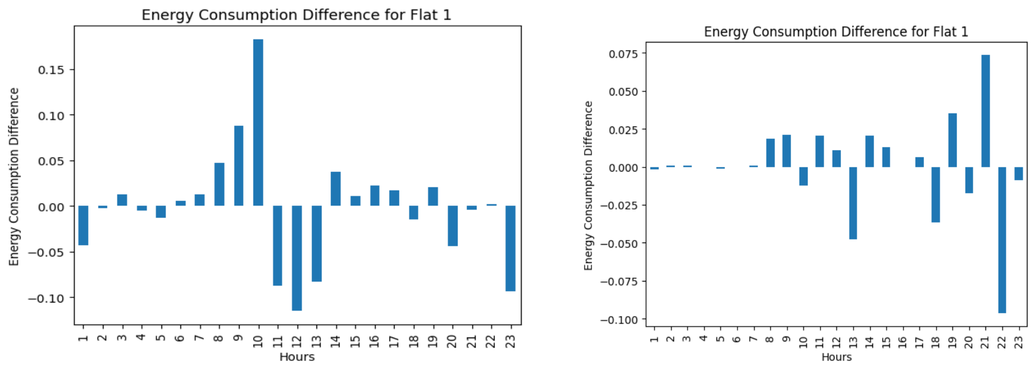


*Figure 29: Differenced Data Plot – a) Winter b) Summer*

*Table 13: ADF Test Results after first order Differencing - Winter*

| | |
|---|---|
| ADF Statistic: | -3.066441 |
| p-value | 0.029138 |
| Critical Value (1%) | -3.809 |
| Critical Value (5%) | -3.022 |
| Critical Value (10%) | -2.651 |

*Table 14: ADF Test Results after first order Differencing - Summer*

| | |
|---|---|
| ADF Statistic: | -28.809440 |
| p-value | 0.000000 |
| Critical Value (1%) | -4.069 |
| Critical Value (5%) | -3.127 |
| Critical Value (10%) | -2.702 |

The test resulted in p-values of 0.029 for winter and 0.0 for summer, which were less than the conventional threshold of 0.05, suggesting the null hypothesis of a unit root could be rejected and the transformed series was stationary. It was then saved for further analysis.

## 6.4 ARIMA Model Configuration

In order to visually understand the correlation structure of the hourly time series data after differencing, Autocorrelation Function (ACF) and Partial Autocorrelation Function (PACF) plots were generated in Python (Figures 30 - 33).

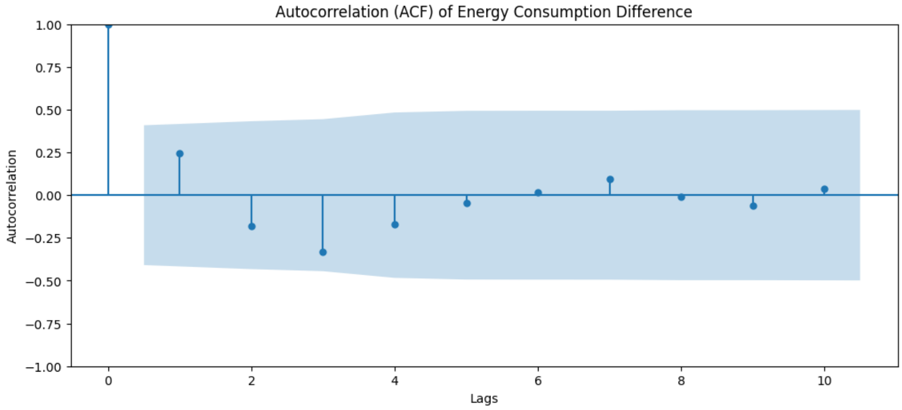


*Figure 30: ACF Plot – Winter dataset*

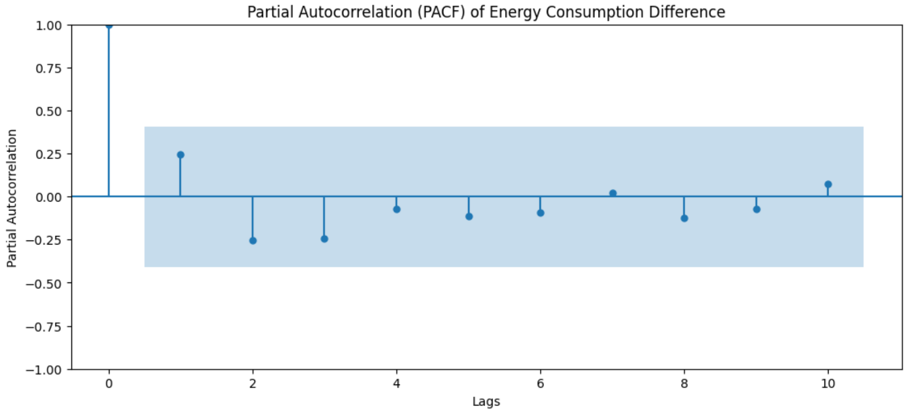


*Figure 31: PACF Plot – Winter dataset*

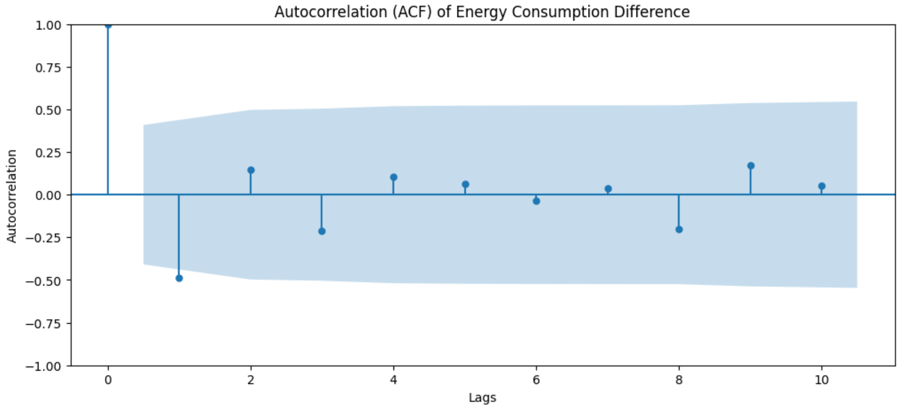


*Figure 32: ACF Plot – Summer dataset*

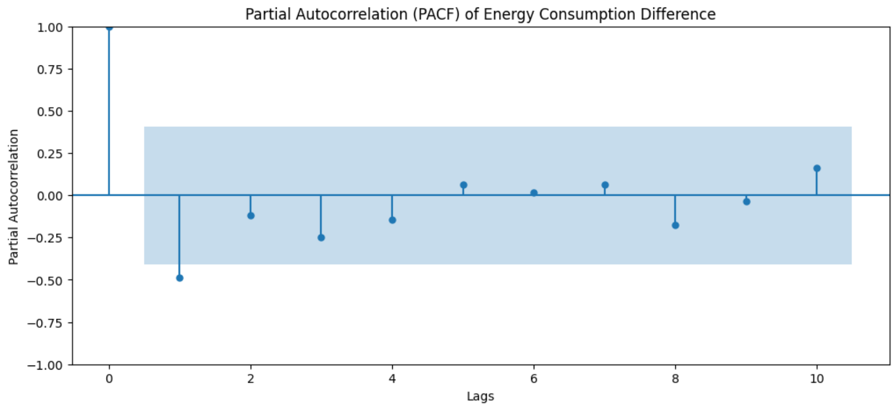


*Figure 33: PACF Plot - Summer dataset*

The significant spikes observed in the ACF and PACF plots presented the initial model parameter estimates. Specifically, the plots suggested starting points of p = 1 for the autoregressive (AR) component and q = 1 for the moving average (MA) component. The differencing value, d, was derived from the previous data transformation step, giving an initial ARIMA order of (1,1,1).

### 6.4.1 Optimal ARIMA Model Parameters

To determine the optimal ARIMA configuration (p, d, q order), various combinations were systematically explored. During this process, two key metrics were used, the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC). Both metrics evaluate model goodness of fit, with a penalty for models that might be too intricate. Essentially, lower AIC or BIC values denote models that strike an ideal balance between accuracy and simplicity.

Using Python, these metrics were generated after model fitting with a goal to select the model showcasing the lowest AIC and BIC scores. The results of the script are decribed in tables 15 and 16, which identified the ARIMA(0, 0, 1) as the optimal order for the hourly winter dataset, and the ARIMA(1,0,0) for the hourly summer dataset. The following sections delve into the evaluation and assessment of how well the chosen ARIMA models predicted hourly energy consumption.

*Table 15: Optimal ARIMA Model Parameters – Winter*

| *Best ARIMA Order* | *(0, 0, 1)* |
|---|---|
| *Best AIC* | -59.23 |
| *Best BIC* | -55.83 |

*Table 16: Optimal ARIMA Model Parameters – Summer*

| *Best ARIMA Order* | *(1, 0, 0)* |
|---|---|
| *Best AIC* | -102.67 |
| *Best BIC* | -99.14 |

## 6.5 Hourly Data Model Evaluations

### 6.5.1 ARIMA(1,1,1) Model

The preliminary ARIMA model, with parameters set to (1,1,1), was employed to predict hourly energy consumption for "Flat 2", using "Flat 1" as the training dataset. The ensuing outcomes, presented both graphically and in tabular form, are displayed below.

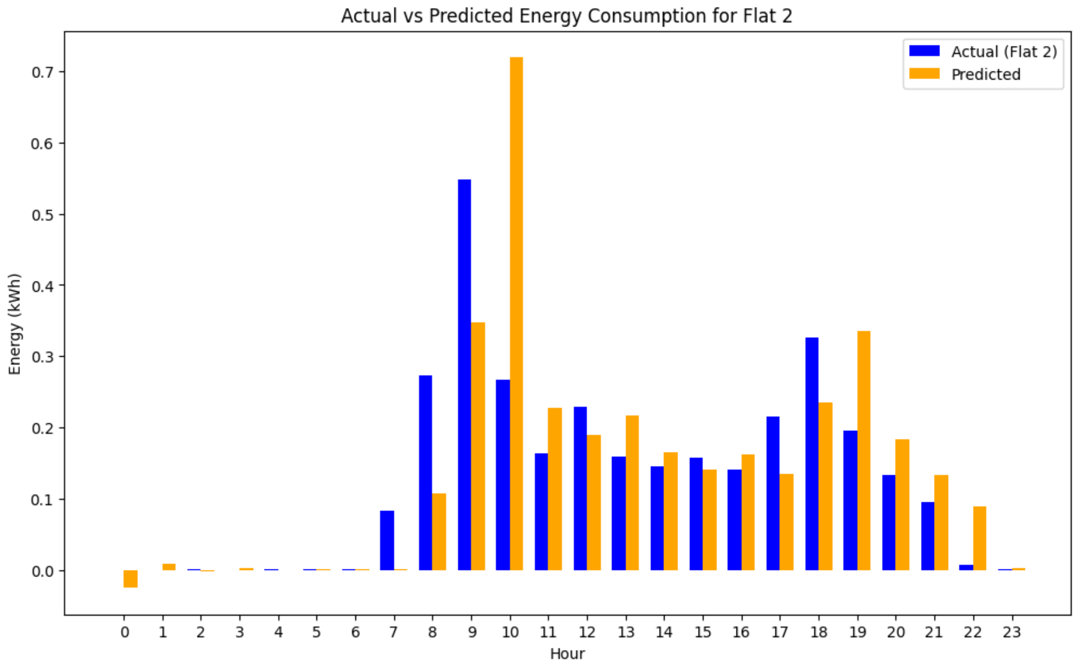


*Figure 34: ARIMA(1,1,1) Results - Winter*

*Table 17: ARIMA(1,1,1) Errors - Winter*

| **Hour** | **Actual** | **Predicted** | **Error** | **Absolute Error** |
|---|---|---|---|---|
| **0** | 0.0000 | -0.0252 | 0.0252 | 0.0252 |
| **1** | 0.0000 | 0.0080 | -0.0080 | 0.0080 |
| **2** | 0.0011 | -0.0026 | 0.0037 | 0.0037 |
| **3** | 0.0000 | 0.0023 | -0.0023 | 0.0023 |
| **4** | 0.0006 | -0.0007 | 0.0013 | 0.0013 |
| **5** | 0.0006 | 0.0010 | -0.0004 | 0.0004 |
| **6** | 0.0006 | 0.0004 | 0.0001 | 0.0001 |
| **7** | 0.0825 | 0.0006 | 0.0819 | 0.0819 |
| **8** | 0.2731 | 0.1069 | 0.1663 | 0.1663 |
| **9** | 0.5479 | 0.3470 | 0.2009 | 0.2009 |
| **10** | 0.2676 | 0.7195 | -0.4520 | 0.4520 |
| **11** | 0.1639 | 0.2273 | -0.0634 | 0.0634 |
| **12** | 0.2285 | 0.1889 | 0.0396 | 0.0396 |
| **13** | 0.1594 | 0.2175 | -0.0581 | 0.0581 |
| **14** | 0.1449 | 0.1656 | -0.0207 | 0.0207 |
| **15** | 0.1577 | 0.1406 | 0.0171 | 0.0171 |
| **16** | 0.1410 | 0.1626 | -0.0216 | 0.0216 |
| **17** | 0.2152 | 0.1345 | 0.0807 | 0.0807 |
| **18** | 0.3255 | 0.2349 | 0.0906 | 0.0906 |
| **19** | 0.1957 | 0.3354 | -0.1398 | 0.1398 |
| **20** | 0.1338 | 0.1838 | -0.0500 | 0.0500 |

| | | | | |
|---|---|---|---|---|
| **21** | 0.0959 | 0.1330 | -0.0371 | 0.0371 |
| **22** | 0.0072 | 0.0889 | -0.0816 | 0.0816 |
| **23** | 0.0006 | 0.0023 | -0.0017 | 0.0017 |
| | | | **RMSE: 0.1181** | **MAE: 0.0685** |

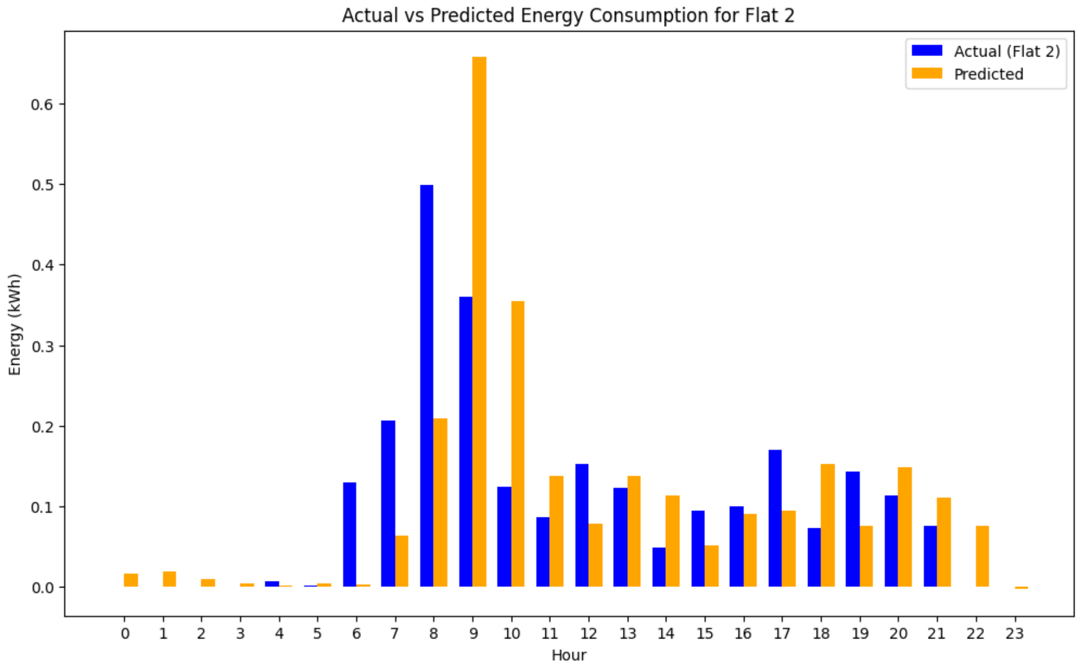


*Figure 35: ARIMA(1,1,1) Results - Summer*

*Table 18: ARIMA(1,1,1) Errors - Summer*

| **Hour** | **Actual** | **Predicted** | **Error** | **Absolute Error** |
|---|---|---|---|---|
| **0** | 0.0000 | 0.0161 | -0.0161 | 0.0161 |
| **1** | 0.0000 | 0.0192 | -0.0192 | 0.0192 |
| **2** | 0.0000 | 0.0092 | -0.0092 | 0.0092 |
| **3** | 0.0000 | 0.0040 | -0.0040 | 0.0040 |
| **4** | 0.0076 | 0.0019 | 0.0057 | 0.0057 |
| **5** | 0.0022 | 0.0042 | -0.0021 | 0.0021 |
| **6** | 0.1292 | 0.0036 | 0.1256 | 0.1256 |
| **7** | 0.2067 | 0.0635 | 0.1432 | 0.1432 |
| **8** | 0.4989 | 0.2085 | 0.2904 | 0.2904 |
| **9** | 0.3599 | 0.6574 | -0.2976 | 0.2976 |
| **10** | 0.1249 | 0.3551 | -0.2302 | 0.2302 |
| **11** | 0.0862 | 0.1376 | -0.0514 | 0.0514 |
| **12** | 0.1521 | 0.0788 | 0.0733 | 0.0733 |
| **13** | 0.1227 | 0.1378 | -0.0152 | 0.0152 |
| **14** | 0.0491 | 0.1141 | -0.0650 | 0.0650 |
| **15** | 0.0949 | 0.0515 | 0.0434 | 0.0434 |
| **16** | 0.1003 | 0.0902 | 0.0101 | 0.0101 |
| **17** | 0.1707 | 0.0949 | 0.0757 | 0.0757 |
| **18** | 0.0731 | 0.1529 | -0.0799 | 0.0799 |
| **19** | 0.1434 | 0.0757 | 0.0677 | 0.0677 |
| **20** | 0.1134 | 0.1485 | -0.0350 | 0.0350 |
| **21** | 0.0763 | 0.1115 | -0.0351 | 0.0351 |

| | | | | |
|---|---|---|---|---|
| **22** | 0.0005 | 0.0758 | -0.0752 | 0.0752 |
| **23** | 0.0000 | -0.0026 | 0.0026 | 0.0026 |
| | | | **RMSE: 0.1122** | **MAE: 0.0739** |

### 6.5.2 ARIMA(0,0,1) and (1,0,0) Models

The optimal ARIMA models (0,0,1) for the winter dataset and (1,0,0) for the summer dataset were also employed to predict hourly energy consumption for "Flat 2", using "Flat 1" as the training dataset. Visual assessments, using bar plots comparing actual consumption values against predicted values, offered a representation of the model's performance. Quantitatively, the model's accuracies were gauged using two metrics: Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE) which were summarised in subsequent tables. These metrics provided insight into the average magnitude of prediction errors and the average absolute differences between predicted and actual values, respectively.

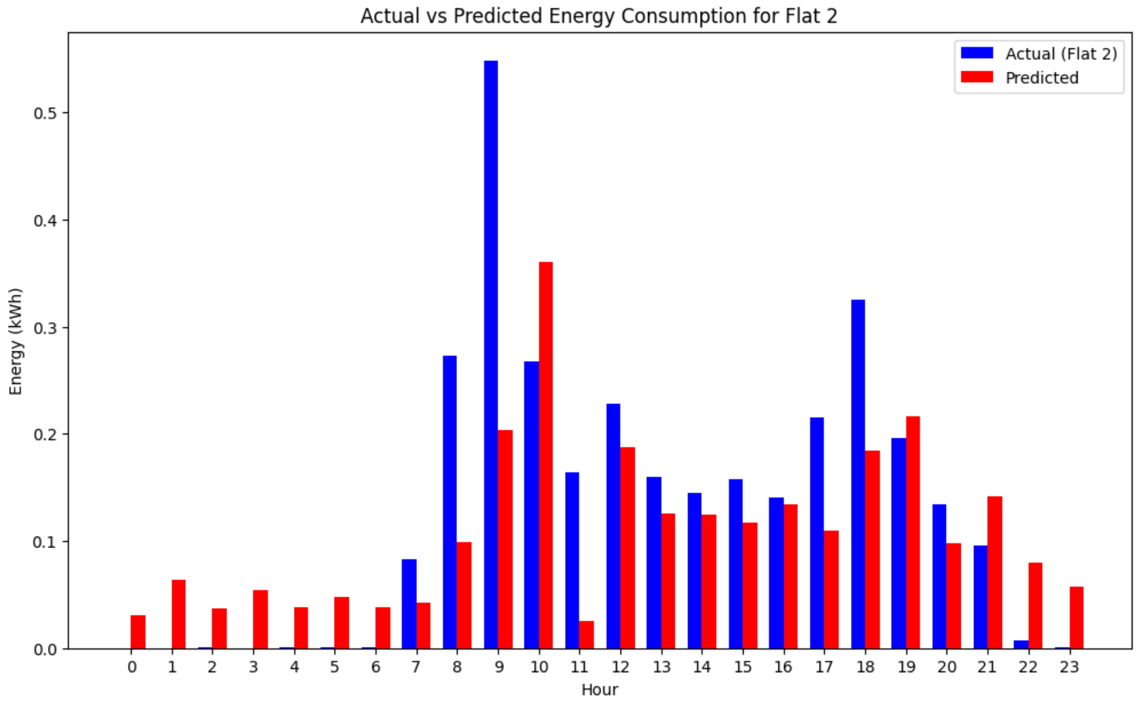


*Figure 36: ARIMA(0,0,1) Results - Winter*

*Table 19: ARIMA(0,0,1) Errors - Winter*

| Hour | Actual | Predicted | Error | Absolute Error |
|---|---|---|---|---|
| **0** | 0.0000 | 0.0303 | -0.0303 | 0.0303 |
| **1** | 0.0000 | 0.0640 | -0.0640 | 0.0640 |
| **2** | 0.0011 | 0.0368 | -0.0357 | 0.0357 |
| **3** | 0.0000 | 0.0540 | -0.0540 | 0.0540 |
| **4** | 0.0006 | 0.0379 | -0.0373 | 0.0373 |
| **5** | 0.0006 | 0.0472 | -0.0467 | 0.0467 |
| **6** | 0.0006 | 0.0378 | -0.0372 | 0.0372 |
| **7** | 0.0825 | 0.0423 | 0.0402 | 0.0402 |
| **8** | 0.2731 | 0.0986 | 0.1746 | 0.1746 |
| **9** | 0.5479 | 0.2040 | 0.3439 | 0.3439 |
| **10** | 0.2676 | 0.3601 | -0.0926 | 0.0926 |

| | | | | |
|---|---|---|---|---|
| **11** | 0.1639 | 0.0248 | 0.1391 | 0.1391 |
| **12** | 0.2285 | 0.1872 | 0.0413 | 0.0413 |
| **13** | 0.1594 | 0.1260 | 0.0334 | 0.0334 |
| **14** | 0.1449 | 0.1249 | 0.0201 | 0.0201 |
| **15** | 0.1577 | 0.1168 | 0.0410 | 0.0410 |
| **16** | 0.1410 | 0.1338 | 0.0072 | 0.0072 |
| **17** | 0.2152 | 0.1099 | 0.1052 | 0.1052 |
| **18** | 0.3255 | 0.1840 | 0.1415 | 0.1415 |
| **19** | 0.1957 | 0.2166 | -0.0210 | 0.0210 |
| **20** | 0.1338 | 0.0982 | 0.0356 | 0.0356 |
| **21** | 0.0959 | 0.1418 | -0.0460 | 0.0460 |
| **22** | 0.0072 | 0.0794 | -0.0721 | 0.0721 |
| **23** | 0.0006 | 0.0570 | -0.0565 | 0.0565 |
| | | | **RMSE: 0.1002** | **MAE: 0.0715** |

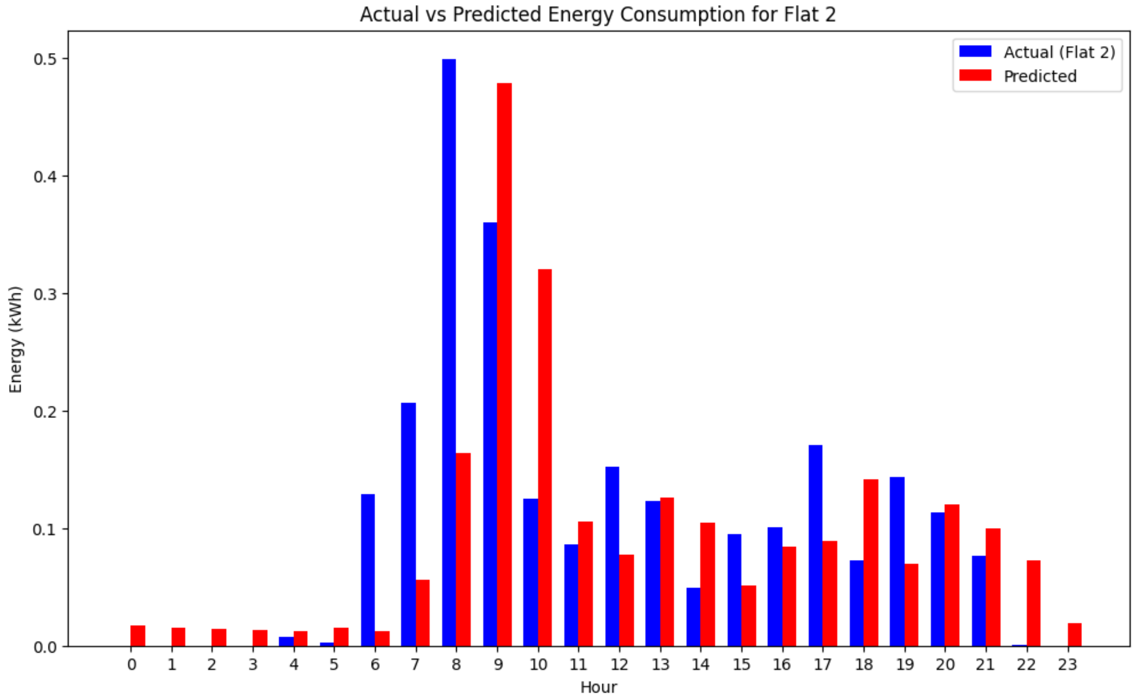


*Figure 37: ARIMA(1,0,0) Results – Summer*

*Table 20: ARIMA(1,0,0) Errors – Summer*

| **Hour** | **Actual** | **Predicted** | **Error** | **Absolute Error** |
|---|---|---|---|---|
| **0** | 0.0000 | 0.0170 | -0.0170 | 0.0170 |
| **1** | 0.0000 | 0.0153 | -0.0153 | 0.0153 |
| **2** | 0.0000 | 0.0142 | -0.0142 | 0.0142 |
| **3** | 0.0000 | 0.0132 | -0.0132 | 0.0132 |
| **4** | 0.0076 | 0.0124 | -0.0048 | 0.0048 |
| **5** | 0.0022 | 0.0152 | -0.0130 | 0.0130 |
| **6** | 0.1292 | 0.0123 | 0.1170 | 0.1170 |
| **7** | 0.2067 | 0.0560 | 0.1507 | 0.1507 |
| **8** | 0.4989 | 0.1640 | 0.3350 | 0.3350 |
| **9** | 0.3599 | 0.4794 | -0.1195 | 0.1195 |
| **10** | 0.1249 | 0.3205 | -0.1956 | 0.1956 |

| | | | | |
|---|---|---|---|---|
| **11** | 0.0862 | 0.1053 | -0.0191 | 0.0191 |
| **12** | 0.1521 | 0.0771 | 0.0751 | 0.0751 |
| **13** | 0.1227 | 0.1263 | -0.0037 | 0.0037 |
| **14** | 0.0491 | 0.1051 | -0.0560 | 0.0560 |
| **15** | 0.0949 | 0.0512 | 0.0437 | 0.0437 |
| **16** | 0.1003 | 0.0846 | 0.0158 | 0.0158 |
| **17** | 0.1707 | 0.0888 | 0.0818 | 0.0818 |
| **18** | 0.0731 | 0.1412 | -0.0682 | 0.0682 |
| **19** | 0.1434 | 0.0698 | 0.0737 | 0.0737 |
| **20** | 0.1134 | 0.1206 | -0.0072 | 0.0072 |
| **21** | 0.0763 | 0.0994 | -0.0230 | 0.0230 |
| **22** | 0.0005 | 0.0729 | -0.0724 | 0.0724 |
| **23** | 0.0000 | 0.0187 | -0.0187 | 0.0187 |
| | | | **RMSE: 0.0993** | **MAE: 0.0647** |

The model's effectiveness in forecasting hourly energy consumption was clear through both visual and quantitative assessments. The bar plots highlighted hours where predicted values closely mirrored actual consumption, and other times when discrepancies arose. With RMSE and MAE metrics lower than the baseline SMA model, the ARIMA models demonstrated superior performance. Summaries of the performance metrics of these models are shown in Tables 21 and 22.

*Table 21: SMA Hourly Forecast Model Performance Summary*

| | ***RMSE*** | | ***MAE*** | |
|---|---|---|---|---|
| | Winter | Summer | Winter | Summer |
| ***SMA*** | 0.1400 | 0.1338 | 0.1083 | 0.0903 |
| ***Enhanced SMA*** | 0.1466 | 0.159 | 0.0901 | 0.107 |

*Table 22: ARIMA Hourly Forecast Model Performance Summary*

| | ***RMSE*** | | ***MAE*** | |
|---|---|---|---|---|
| | Winter | Summer | Winter | Summer |
| ***ARIMA(1,1,1)*** | 0.1181 | 0.1122 | 0.0685 | 0.0739 |
| ***ARIMA(0,0,1)*** | 0.1002 | | 0.0715 | |
| ***ARIMA(1,0,0)*** | | 0.0993 | | 0.0647 |

Although the ARIMA models exhibited better performance, the graphs in figures 34 to 37 show an initial underprediction during the morning demand, which is then offset by an over-prediction two hours later. Given that heating systems predominantly focus on the morning peak demand, it may be beneficial to adjust the timestamps from the ARIMA outputs to achieve a more accurate fit. This modification is suggested for consideration in future work to enhance the model.

## 6.6 Daily Data Model Evaluations

### 6.6.1 Winter Daily Models

The previously described methods were used to calculate the best ARIMA order for the daily winter dataset and the results together with the corresponding ACF and PACF plots are shown below:

Table 23: Optimal ARIMA parameters - Winter

| *Best ARIMA Order* | *(0, 0, 4)* |
|---|---|
| *Best AIC* | 529.17 |
| *Best BIC* | 544.55 |

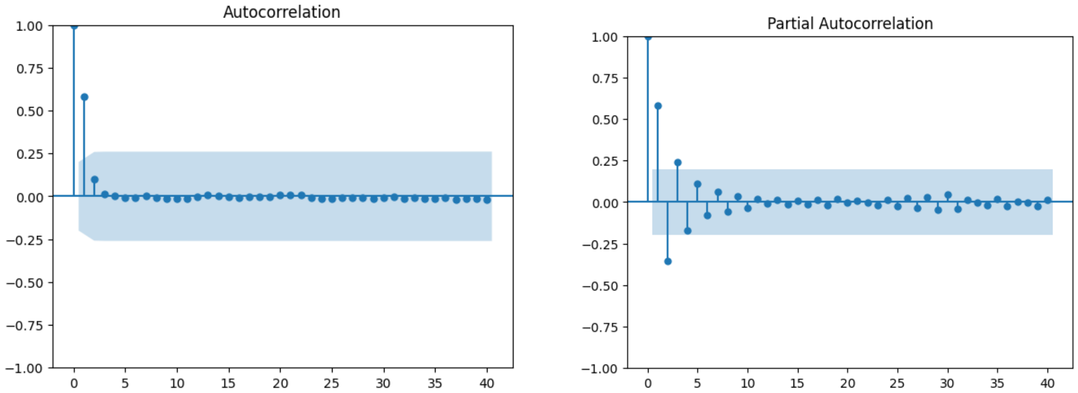


Figure 38: ACF and PACF Plots for Daily Usage – Winter

The ARIMA(0,0,4) Model for the winter dataset was fit into the script and the results are as follows:

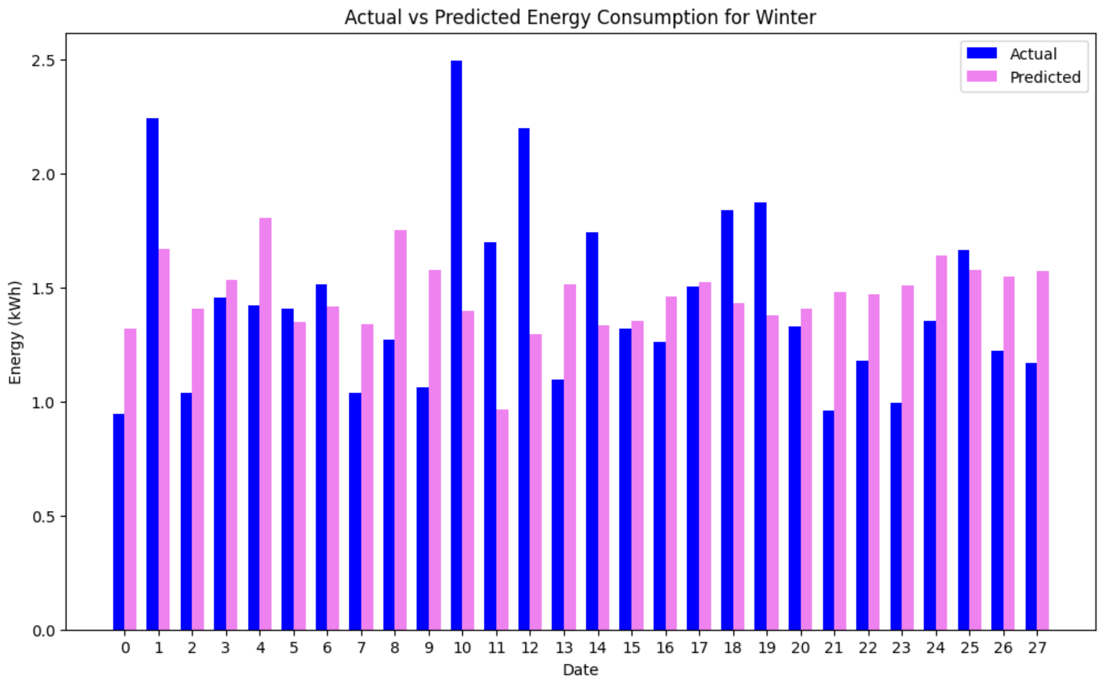


Figure 39: ARIMA(0,0,4) Results – Winter

Table 24: ARIMA(0,0,4) Evaluation Metrics

| *ARIMA Order* | *(0, 0, 4)* |
|---|---|
| *RMSE* | 0.4503 |
| *MAE* | 0.3727 |

Despite the ARIMA(0,0,4) being the optimal order from the AIC and BIC tests, another model, ARIMA(0,0,6) was fit based on trial and error to see if better results would be observed. The results are as follows:

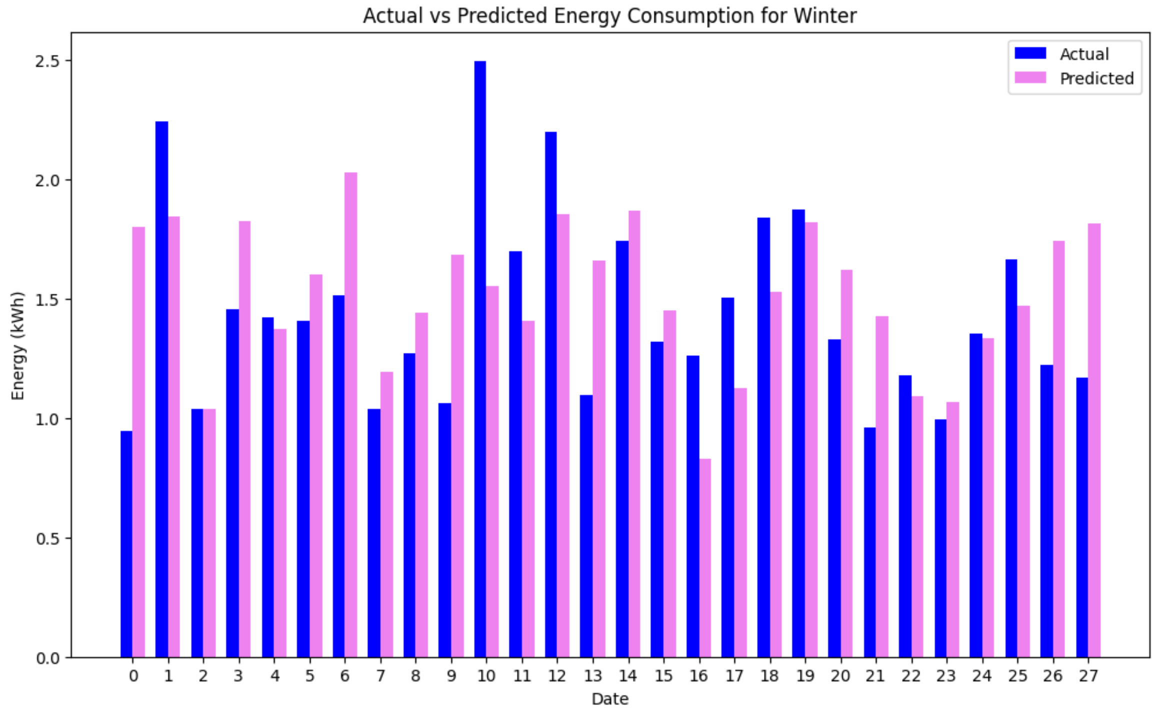


*Figure 40: ARIMA(0,0,6) Results – Winter*

*Table 25: ARIMA(0,0,6) Evaluation Metrics - Winter*

| *ARIMA Order* | *(0, 0, 6)* |
|---:|---|
| *RMSE* | 0.4089 |
| *MAE* | 0.3285 |

### 6.6.2 Summer Daily Models

The AIC and BIC tests were carried out on the daily datasets for the summer season with the results as follows in addition to the corresponding ACF and PACF plots.

| *Best ARIMA Order* | *(4, 1, 0)* |
|---:|---|
| *Best AIC* | 517.97 |
| *Best BIC* | 530.84 |

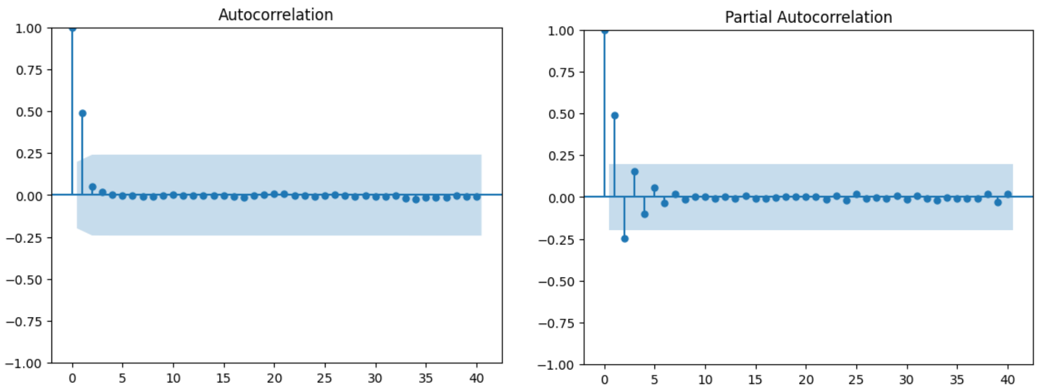


*Figure 41: ACF and PACF Plots for Daily Usage – Summer*

The ARIMA(4,1,0) Model for the summer dataset was fit into the script and the results are as follows:

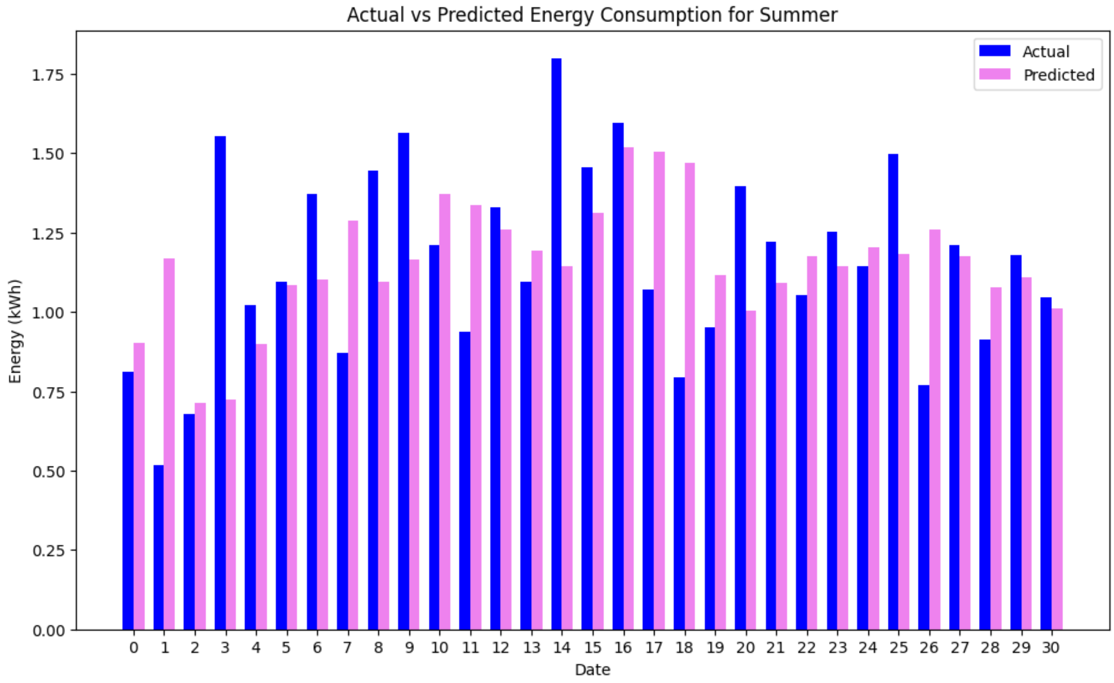


Figure 42: ARIMA(4,1,0) Results – Summer

| *ARIMA Order* | *(4, 1, 0)* |
|---|---|
| *RMSE* | 0.3395 |
| *MAE* | 0.2574 |

Similar to the winter dataset, despite the ARIMA(4,1,0) being the optimal order from the AIC and BIC tests, another model, ARIMA(1,0,1) was fit based on trial and error to see if better results would be observed. The results are as follows:

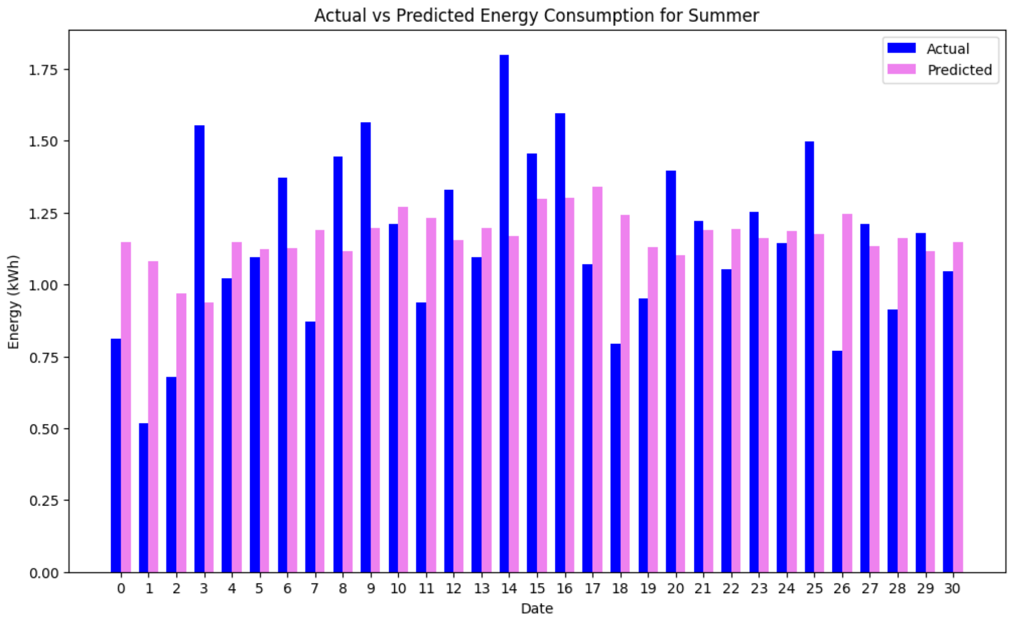


Figure 43: ARIMA(1,0,1) Results - Summer

Table 26: ARIMA(1,0,1) Evaluation Metrices - Summer

| *ARIMA Order* | *(1, 0, 1)* |
|---|---|
| *RMSE* | 0.3003 |
| *MAE* | 0.2490 |

Table 27 displays the performance metric summaries for the models that predicted daily energy consumption. While the baseline SMA model had better RMSE and MAE metrics compared to the ARIMA models, the ARIMA models provided superior visual representations of the predicted values. These insights suggest that more in-depth research and analysis are required to optimise the models for the most accurate predictions.

Table 27: Daily Forecast Model Performance Summary

| | ***RMSE*** | | ***MAE*** | |
|---|---|---|---|---|
| | Winter | Summer | Winter | Summer |
| ***SMA*** | 0.3759 | 0.2662 | 0.2864 | 0.2169 |
| ***ARIMA(0,0,4)*** | 0.4503 | | 0.3727 | |
| ***ARIMA(0,0,6)*** | 0.4089 | | 0.3285 | |
| ***ARIMA(4,1,0)*** | | 0.3395 | | 0.2574 |
| ***ARIMA(1,0,1)*** | | 0.3003 | | 0.2490 |

## 6.7 Model Evaluation Using Residual Errors

### 6.7.1 Hourly Data Residual Errors

Evaluation of the hourly ARIMA forecasting models for energy consumption revealed promising fits and plots of their residual errors are shown below. The trendless nature of the plots suggested that most systematic information from the original series had been captured by the models. With mean residual errors of near zero (0.022 for Winter and 0.010 for Summer), the models did not exhibit persistent biases further attesting to their efficacies. The bell-shaped density plots centred around zero highlighted the normal distribution-like behaviour of residuals, which is a favourable attribute.

The histograms displayed a leftward skew, suggesting a minor asymmetry in the error distribution. Meanwhile, the ACF plots revealed limited residual autocorrelation, signifying that the model effectively accounted for the majority of the series' inherent dynamics.

Winter mean residual error: 0.022

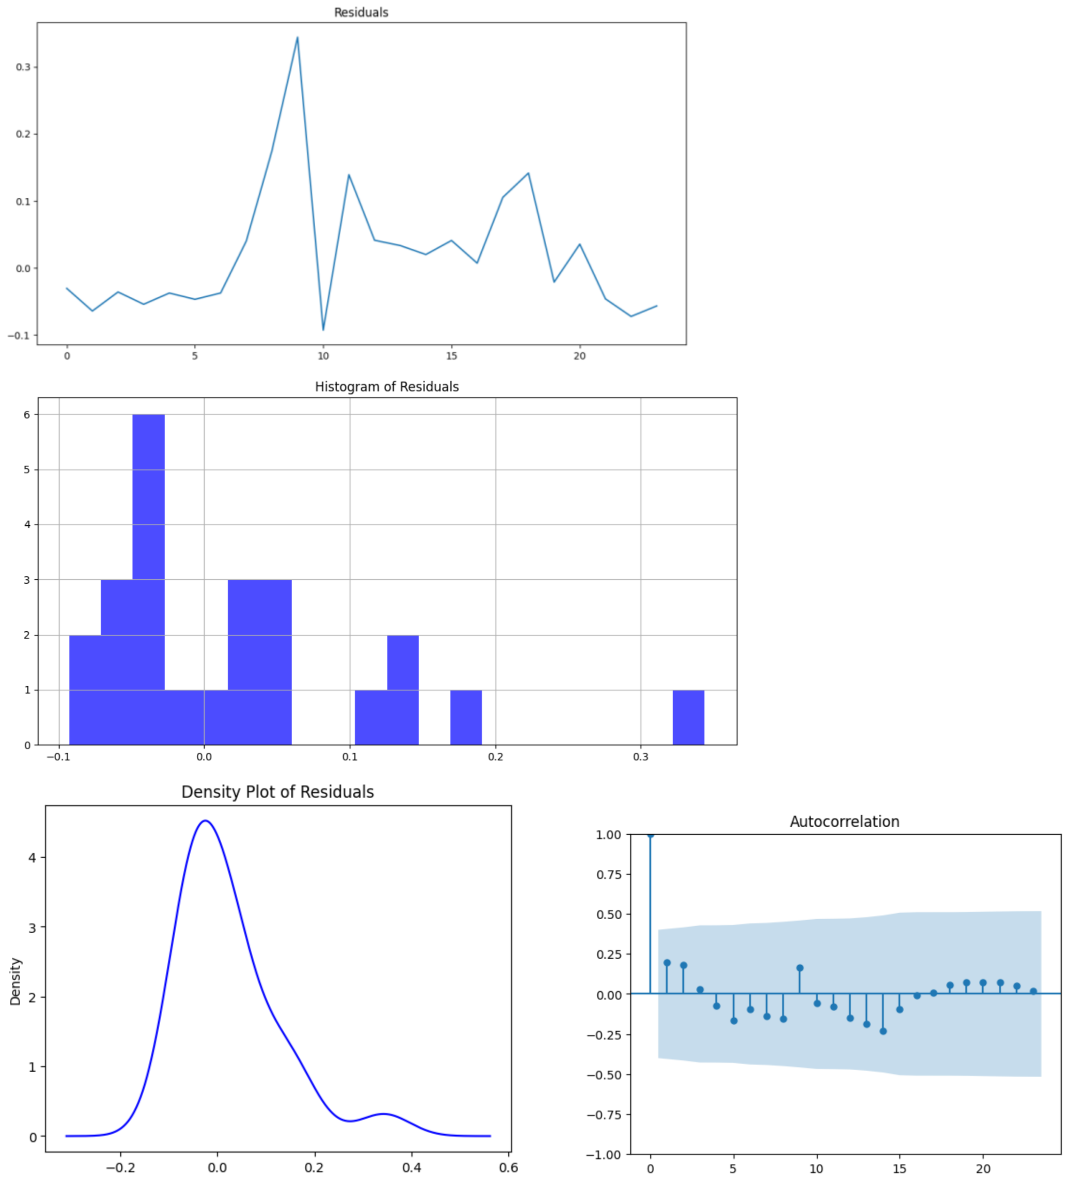


*Figure 44: Hourly Forecast Residual Errors – Winter*

Summer mean residual error: 0.010

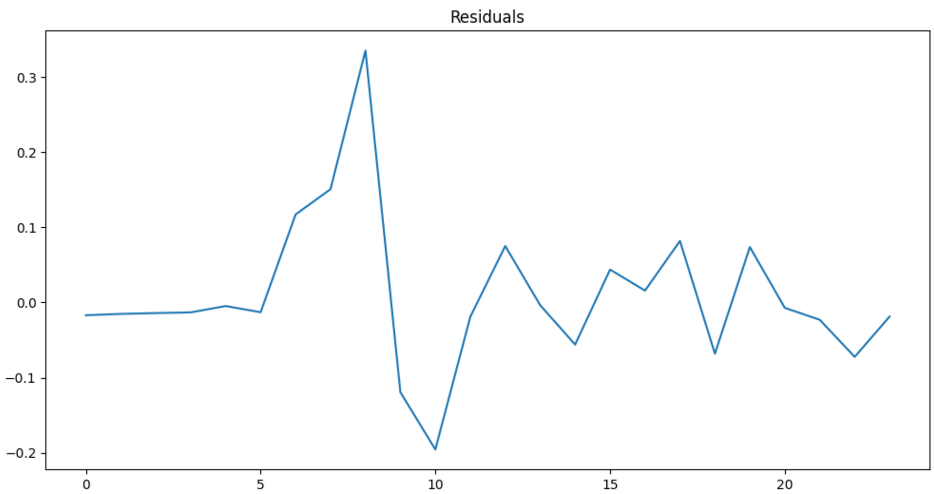

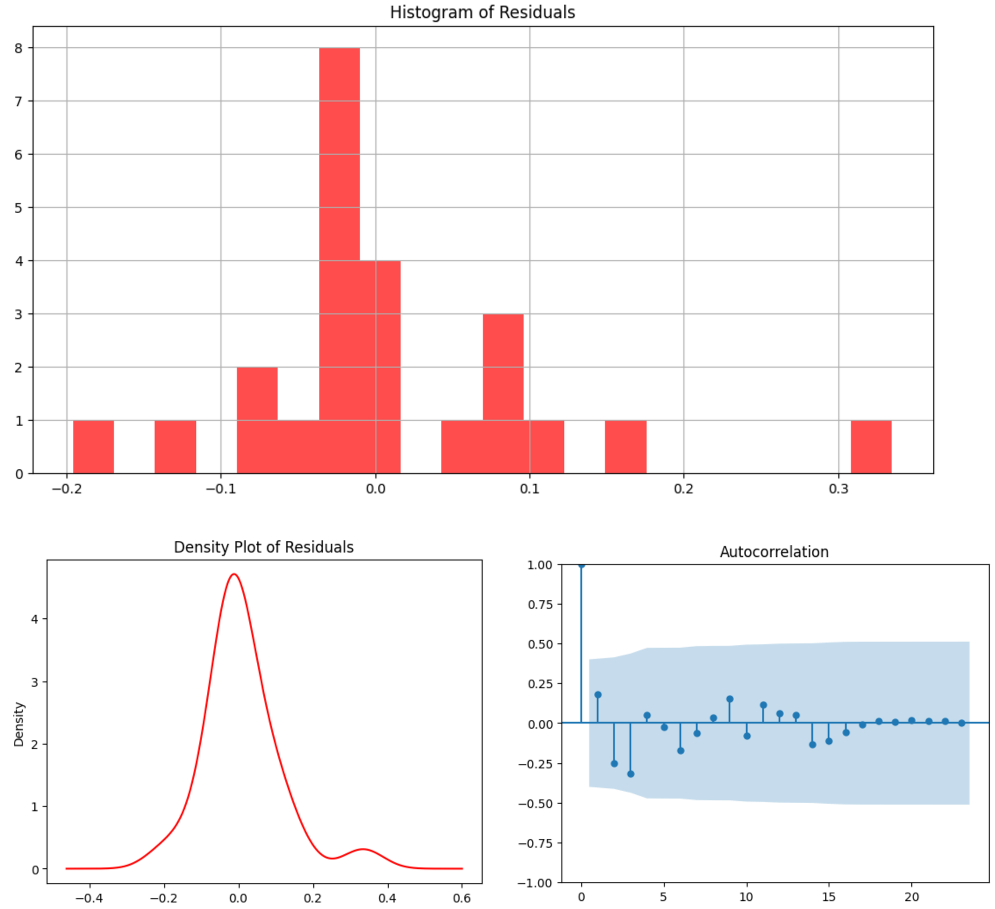


*Figure 45: Hourly Forecast Residual Errors – Summer*

### 6.7.2 Daily Data Residual Errors

Similar observations were seen with the residual error plots for the daily forecasting models. The plots displayed near-zero mean errors, signifying negligible bias while the trendless nature of these plots indicated that most systematic information from the original series was captured by the models. The histograms exhibited skewness as well, indicating some asymmetry in the error distribution. Additionally, the ACF plots showed little autocorrelation in the residuals, indicating that the models effectively captured the series' main patterns.

Winter mean residual error: -0.0778

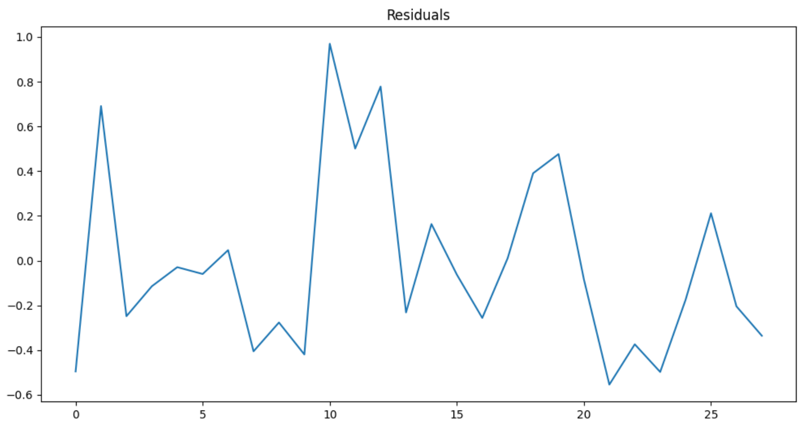

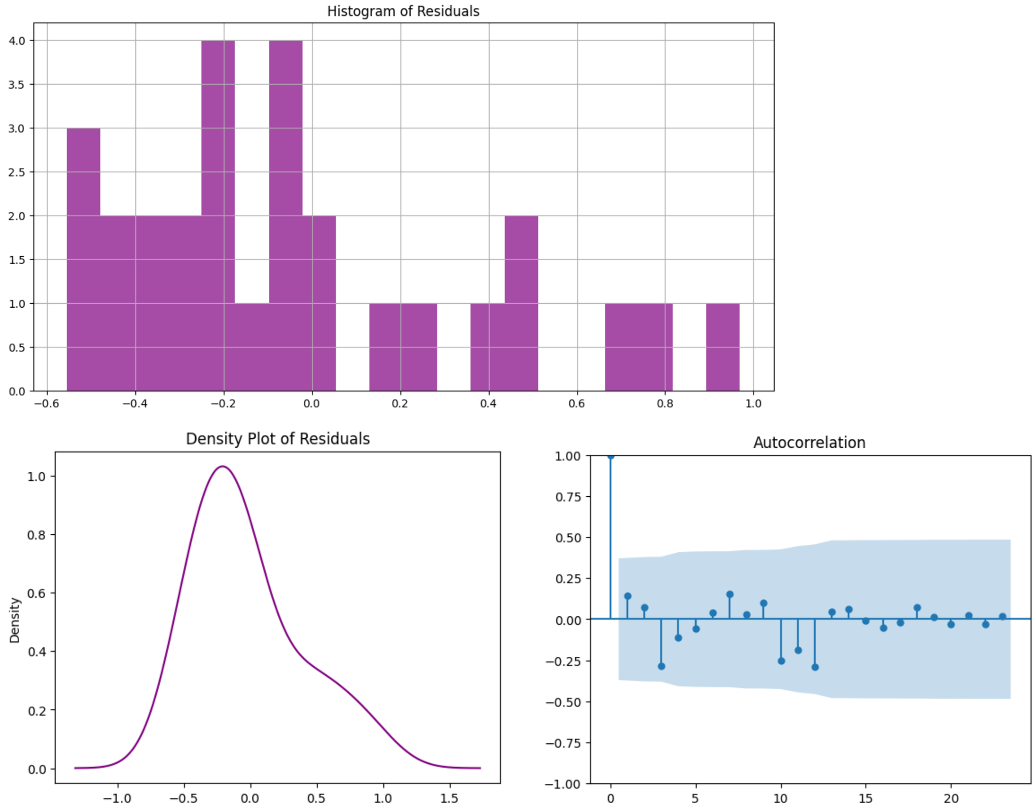


*Figure 46: Daily Forecast Residual Errors - Winter*

Summer mean residual error: 0.0105

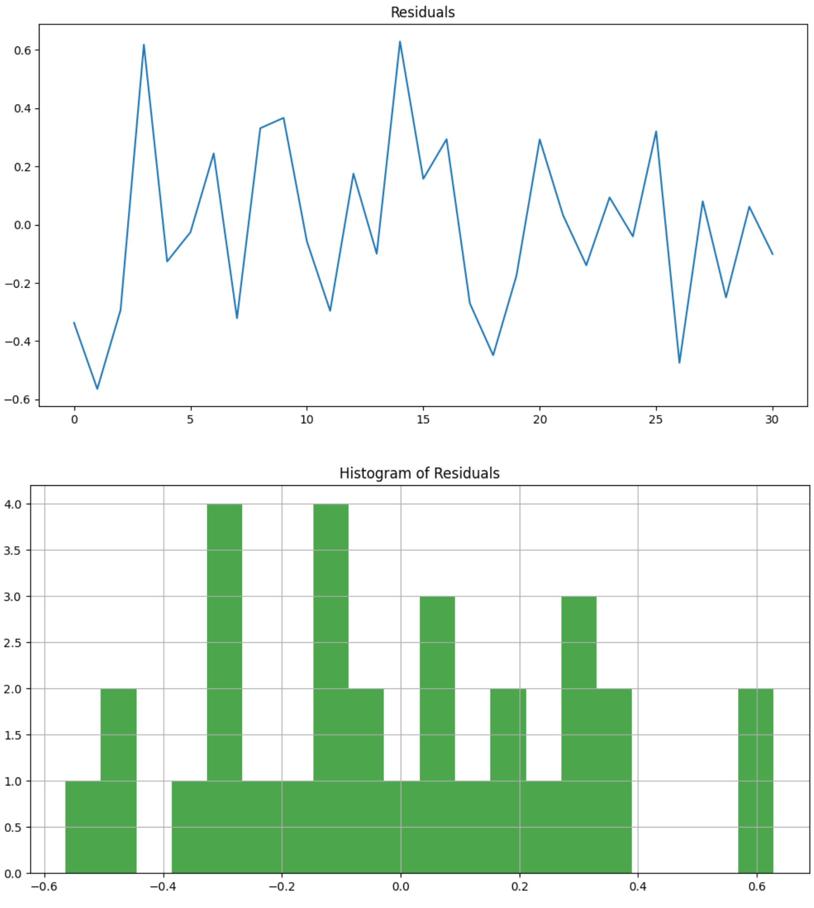

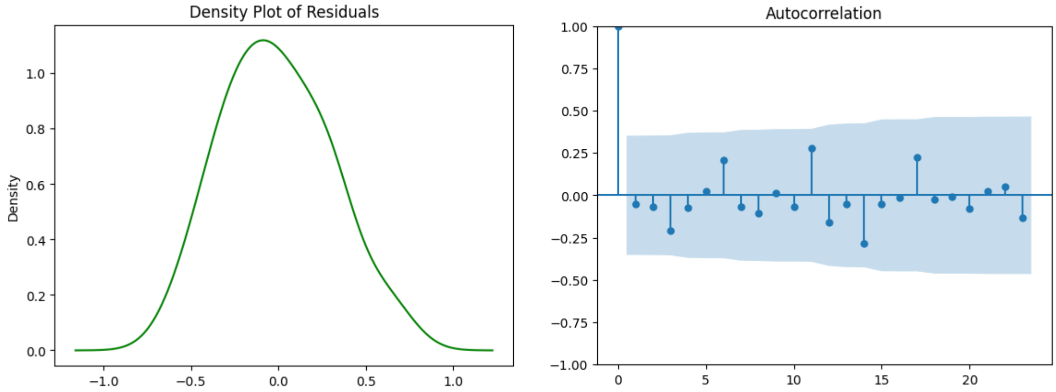


*Figure 47: Daily Forecast Residual Errors - Summer*

# 7.0 Discussion and Conclusions

Time series analysis of Domestic Hot Water (DHW) energy consumption was performed on real-world data from the West Whins residential building in Findhorn. Various techniques were assessed, and the ARIMA method was selected to construct DHW consumption forecasting models for both hourly and daily consumptions, comparing their performances in winter and summer time periods.

A Simple Moving Average (SMA) model was developed and evaluated on both the hourly and daily consumption datasets for benchmarking. This basic model served as a reference point for comparing more advanced models using the Root Mean Squared Error (RMSE) and Mean Squared Error (MAE) performance metrics. Before developing the ARIMA models, the datasets were tested for stationarity using Augmented Dickey-Fuller (ADF) tests. The results of the hourly data highlighted non-stationary characteristics, leading to performance of a first order differentiation to achieve stationarity.

To identify the parameters for the ARIMA model, Autocorrelation Function (ACF) and Partial Autocorrelation Function (PACF) plots of the datasets were generated. Notable spikes observed on the plots suggested the initial ‘p’ and ‘q’ values to be used for the ARIMA modelling, with the ‘d’ value based on the stationarity of the data. To further pinpoint the optimal parameters for the best fit model, the Alkaline Information Criterion (AIC) and Bayesian Information Criterion (BIC) methods were employed. The results from these tests gave the optimal ARIMA models for the hourly and daily consumption datasets.

Upon examining the predictive capabilities of the hourly forecasting models, it was observed that the ARIMA (1,1,1); (0,0,1) and (1,0,0) models outperformed the baseline SMA model, as indicated by their lower RMSE and MAE values. For the winter dataset, the best-performing model was ARIMA(0,0,1) with an RMSE of 0.1002 and an MAE of 0.0715. Conversely, for the summer dataset, the ARIMA(1,0,0) model stood out, with an RMSE of 0.0993 and an MAE of 0.0647.

The residual errors were also plotted for further analysis of the models. For both hourly and daily datasets, these errors were close to zero, suggesting minimal biases. The absence of trends in these plots suggested that the models effectively captured most of the systematic information from the original series. However, histograms displayed skewness, suggesting some unevenness in the error distribution.

Despite the low RMSE and MAE values, plots comparing the actual versus predicted values of hourly energy demand revealed a noticeable underprediction during the morning peak times, followed by a subsequent overprediction two hours later. It is therefore proposed that the timestamps of the outputs of the models be adjusted for more accurate model fitting in future work.

A number of limitations arose during this study, including time constraints, insufficient data details, and a skillset gap in Python programming language that restricted exploration into more intricate models. Nevertheless, the ARIMA models presented in this study offer a promising foundation for incorporation into the West Whins energy centre's intelligent control system in Findhorn. Further validation and comparisons with other forecasting methods could also help refine the model.

# 7.0 References


[1] U. Nations. "Five ways to jump-start the renewable energy transition now." www.un.org/en/climatechange/raising-ambition/renewable-energy-transition (accessed 2023).

[2] E. Department of Business, and Industrial Strategy (BEIS), "Updated energy and emissions projections 2021 to 2040," 2021. [Online]. Available: Energy and emissions projections - GOV.UK (www.gov.uk)

[3] C. E. P. R. I. (CEPRI), "The White Paper of Chinese Steel Industry Electrification," Internal communication with China Electric Power Research Institute, 2020.

[4] I. R. E. A. (IRENA), "INNOVATION LANDSCAPE FOR SMART ELECTRIFICATION - Decarbonising End-Use Sectors with Renewable Power," 2023. [Online]. Available: https://www.irena.org/Publications/2023/Jun/Innovation-landscape-for-smart-electrification

[5] H. Averfalk, P. Ingvarsson, U. Persson, M. Gong, and S. Werner, "Large heat pumps in Swedish district heating systems," *Renewable and Sustainable Energy Reviews,* vol. 79, pp. 1275-1284, 2017/11/01/ 2017, doi: https://doi.org/10.1016/j.rser.2017.05.135.

[6] S. Madeddu *et al.*, "The CO2 reduction potential for the European industry via direct electrification of heat supply (power-to-heat)," *Environmental Research Letters,* vol. 15, no. 12, p. 124004, 2020/11/25 2020, doi: 10.1088/1748-9326/abbd02.

[7] I. R. E. A. (IRENA), "IRENA (2020a), Global Renewables Outlook: Energy Transformation 2050," ed. Abu Dhabi, 2020.

[8] A. S. Tanenbaum and M. van Steen, *Distributed Systems*. CreateSpace Independent Publishing Platform, 2017.

[9] T. Cuvelier, "Embedding reservoirs in industrial models to exploit their flexibility," *SN Applied Sciences,* vol. 2, 12/01 2020, doi: 10.1007/s42452-020-03925-2.

[10] A. G. Swisscom Energy Solutions. "tiKo-using heat pumps and other residential devices to support integration of renewable energies." www.swisscom.ch/en/about/news/2019/03/14-tiko-international.html (accessed.

[11] H2GreenSteel. "How digitalization of hydrogen production will enable better, cleaner industries." (accessed.

[12] R. Petrichenko, K. Baltputnis, A. Sauhats, and D. Sobolevsky, "District heating demand short-term forecasting," in *2017 IEEE International Conference on Environment and Electrical Engineering and 2017 IEEE Industrial and Commercial Power Systems Europe (EEEIC / I&CPS Europe)*, 6-9 June 2017 2017, pp. 1-5, doi: 10.1109/EEEIC.2017.7977633.

[13] F. Bünning, A. Bollinger, P. Heer, R. Smith, and J. Lygeros, *Improved day ahead heating demand forecasting by online correction methods*. 2019.

[14] E. Saloux and J. A. Candanedo, "Forecasting District Heating Demand using Machine Learning Algorithms," *Energy Procedia,* vol. 149, pp. 59-68, 2018/09/01/ 2018, doi: https://doi.org/10.1016/j.egypro.2018.08.169.

[15] I. R. E. A. (IRENA), "Innovation landscape brief: Artificial intelligence and big data," Abu Dhabi, 2019. [Online]. Available: https://www.irena.org/publications/2019/Sep/Artificial-Intelligence-and-Big-Data?ris=1

[16] K. Förderer, M. Ahrens, K. Bao, I. Mauser, and H. Schmeck, "Modeling flexibility using artificial neural networks," *Energy Informatics,* vol. 1, 10/10 2018, doi: 10.1186/s42162-018-0024-4.

[17] G. E. P. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, *Time series analysis: forecasting and control*. John Wiley & Sons, 2015.

[18] A. S. Ahmad *et al.*, "A review on applications of ANN and SVM for building electrical energy consumption forecasting," *Renewable and Sustainable Energy Reviews,* vol. 33, pp. 102–109, 05/01 2014, doi: 10.1016/j.rser.2014.01.069.

[19] S. Barja-Martinez, M. Aragüés-Peñalba, I. Munné, P. Lloret-Gallego, E. Bullich, and R. Villafafila-Robles, "Artificial intelligence techniques for enabling Big Data services in distribution networks: A review," *Renewable and Sustainable Energy Reviews,* vol. 150, p. 111459, 10/01 2021, doi: 10.1016/j.rser.2021.111459.

[20] F. Ecovillage. "Findhorn Ecovillage, New Frontiers for Sustainability." (accessed 2023).

[21] A. L. Graeme Flett, Paul Tuohy, "Design of an Integrated Smart Energy System (SES) Service for Sustainable Buildings and Districts," ESRU ETP Report, Glasgow, 2020.

[22] R. Saidur, H. H. Masjuki, and M. Y. Jamaluddin, "An application of energy and exergy analysis in residential sector of Malaysia," *Energy policy,* vol. 35, no. 2, pp. 1050-1063, 2007.

[23] R. Yao and K. Steemers, "A method of formulating energy load profile for domestic buildings in the UK," *Energy and buildings,* vol. 37, no. 6, pp. 663-671, 2005.

[24] P. Geudens, "Water supply statistics 2007," *Fivewin, Rijswijk,* vol. 83, p. 6259, 2008.

[25] L. Gelažanskas and K. A. A. Gamage, "Forecasting Hot Water Consumption in Residential Houses," *Energies,* vol. 8, no. 11, pp. 12702-12717doi: 10.3390/en81112336.

[26] W. Velicer and J. Fava, "Time Series Analysis," vol. 2, 2003.

[27] J.-S. Chou and N.-T. Ngo, "Time series analytics using sliding window metaheuristic optimization-based machine learning system for identifying building energy consumption patterns," *Applied energy,* vol. 177, pp. 751-770, 2016.

[28] J.-S. Chou and D.-S. Tran, "Forecasting energy consumption time series using machine learning techniques based on usage patterns of residential householders," *Energy,* vol. 165, pp. 709-726, 2018/12/15/ 2018, doi: https://doi.org/10.1016/j.energy.2018.09.144.

[29] J. G. De Gooijer and R. J. Hyndman, "25 years of time series forecasting," *International Journal of Forecasting,* vol. 22, no. 3, pp. 443-473, 2006/01/01/ 2006, doi: https://doi.org/10.1016/j.ijforecast.2006.01.001.

[30] A. S. Bharatpur, *A LITERATURE REVIEW ON TIME SERIES FORECASTING METHODS*. 2022.

[31] J. Brownlee, *Introduction to Time Series Forecasting With Python*, v1.12 ed. 2017, p. 360.

[32] R. J. Hyndman and Y. Khandakar, "Automatic Time Series Forecasting: The forecast Package for R," *Journal of Statistical Software,* vol. 27, no. 3, pp. 1 - 22, 07/29 2008, doi: 10.18637/jss.v027.i03.

[33] C. Deb, F. Zhang, J. Yang, S. E. Lee, and K. W. Shah, "A review on time series forecasting techniques for building energy consumption," *Renewable and Sustainable Energy Reviews,* vol. 74, pp. 902-924, 2017/07/01/ 2017, doi: https://doi.org/10.1016/j.rser.2017.02.085.

[34] A. Delorme-Costil and J. J. Bezian, "Forecasting Domestic Hot Water Demand in Residential House Using Artificial Neural Networks," in *2017 16th IEEE International Conference on Machine Learning and Applications (ICMLA)*, 18-21 Dec. 2017 2017, pp. 467-472, doi: 10.1109/ICMLA.2017.0-117.

[35] P. Compagnon, A. Lomet, M. Reyboz, and M. Mermillod, "Domestic Hot Water Forecasting for Individual Housing with Deep Learning," 12/30 2022.

[36] A. Bauer, M. Züfle, N. Herbst, and S. Kounev, *Best Practices for Time Series Forecasting (Tutorial Paper)*. 2019.

[37] M. Sheha, "Using Real-Time Electricity Prices to Leverage Electrical Energy Storage and Flexible Loads in a Smart Grid Environment Utilizing Machine Learning Techniques," (in English), *Processes,* vol. 7, no. 12, p. 870, 2019 2021-07-12 2019, doi: https://doi.org/10.3390/pr7120870.

[38] R. Zrieq *et al.*, "Time-Series Analysis and Healthcare Implications of COVID-19 Pandemic in Saudi Arabia," *Healthcare*, vol. 10, no. 10, doi: 10.3390/healthcare10101874.

[39] K. Naik and A. Ianakiev, "Heat demand prediction: A real-life data model vs simulated data model comparison," *Energy Reports,* vol. 7, pp. 380-388, 2021/10/01/ 2021, doi: https://doi.org/10.1016/j.egyr.2021.08.093.

[40] I. Khandelwal, R. Adhikari, and G. Verma, "Time series forecasting using hybrid ARIMA and ANN models based on DWT decomposition," *Procedia Computer Science,* vol. 48, pp. 173-179, 2015.

[41] G. P. Zhang, "Time series forecasting using a hybrid ARIMA and neural network model," *Neurocomputing,* vol. 50, pp. 159-175, 2003/01/01/ 2003, doi: https://doi.org/10.1016/S0925-2312(01)00702-0.

[42] A. O. Anele, E. Todini, Y. Hamam, and A. M. Abu-Mahfouz, "Predictive Uncertainty Estimation in Water Demand Forecasting Using the Model Conditional Processor," *Water*, vol. 10, no. 4, doi: 10.3390/w10040475.

[43] V. Vámos and M. Horváth, "Residential DHW consumption analysis for multifamily buildings supplied by district heating," *Thermal Science,* vol. 26, pp. 190-190, 01/01 2021, doi: 10.2298/TSCI200505190V.

[44] J. C. Evarts and L. G. Swan, "Domestic hot water consumption estimates for solar thermal system sizing," *Energy and Buildings,* vol. 58, pp. 58-65, 2013.

[45] K. Ahmed, P. Pylsy, and J. Kurnitski, "Monthly domestic hot water profiles for energy calculation in Finnish apartment buildings," *Energy and Buildings,* vol. 97, 06/15 2015, doi: 10.1016/j.enbuild.2015.03.051.

[46] D. H. Wolpert and W. G. Macready, "No free lunch theorems for optimization," *IEEE Transactions on Evolutionary Computation*, vol. 1, no. 1, pp. 67-82, 1997, doi: 10.1109/4235.585893.

[47] J. Fattah, L. Ezzine, Z. Aman, H. Moussami, and A. Lachhab, "Forecasting of demand using ARIMA model," *International Journal of Engineering Business Management*, vol. 10, p. 184797901880867, 10/29 2018, doi: 10.1177/1847979018808673.

[48] G. E. P. Box and G. M. Jenkins, *Time Series Analysis: Forecasting and Control*, 1 ed. San Francisco: Holden-Day, 1970.

[49] I. Msiza, F. Nelwamondo, and T. Marwala, "Water Demand Prediction using Artificial Neural Networks and Support Vector Regression," *Journal of Computers*, vol. 3, 11/01 2008, doi: 10.4304/jcp.3.11.1-8.

[50] A. Candelieri *et al.*, "Tuning hyperparameters of a SVM-based water demand forecasting system through parallel global optimization," *Computers & Operations Research*, vol. 106, pp. 202-209, 2019/06/01/ 2019, doi: https://doi.org/10.1016/j.cor.2018.01.013.

[51] L. Mu, F. Zheng, R. Tao, Z. Qingzhou, and Z. Kapelan, "Hourly and Daily Urban Water Demand Predictions Using a Long Short-Term Memory Based Model," *Journal of Water Resources Planning and Management*, vol. 146, p. 05020017, 09/01 2020, doi: 10.1061/(ASCE)WR.1943-5452.0001276.

[52] P. I. Karamaziotis, A. Raptis, K. Nikolopoulos, K. Litsiou, and V. Assimakopoulos, "An empirical investigation of water consumption forecasting methods," *International Journal of Forecasting*, vol. 36, no. 2, pp. 588-606, 2020/04/01/ 2020, doi: https://doi.org/10.1016/j.ijforecast.2019.07.009.

[53] L.-G. Maltais and L. Gosselin, "Predictability analysis of domestic hot water consumption with neural networks: From single units to large residential buildings," *Energy*, vol. 229, 2021, doi: 10.1016/j.energy.2021.120658.

[54] L. Gelazanskas and K. Gamage, *Forecasting Hot Water Consumption in Dwellings Using Artificial Neural Networks*. 2015.

[55] A. Jain, A. Varshney, and U. Joshi, "Short-Term Water Demand Forecast Modelling at IIT Kanpur Using Artificial Neural Networks," *Water Resources Management*, vol. 15, pp. 299-321, 10/01 2001, doi: 10.1023/A:1014415503476.

[56] A. Anele, Y. Hamam, A. Abu-Mahfouz, and E. Todini, "Overview, Comparative Assessment and Recommendations of Forecasting Models for Short-Term Water Demand Prediction," *Water*, vol. 9, p. 887, 11/13 2017, doi: 10.3390/w9110887.

[57] M. Romano and Z. Kapelan, "Adaptive water demand forecasting for near real-time management of smart water distribution systems," *Environmental Modelling & Software*, vol. 60, pp. 265-276, 2014/10/01/ 2014, doi: https://doi.org/10.1016/j.envsoft.2014.06.016.

[58] G. Chen, T. Long, J. Xiong, and Y. Bai, "Multiple Random Forests Modelling for Urban Water Consumption Forecasting," *Water Resources Management*, vol. 31, no. 15, pp. 4715-4729, 2017/12/01 2017, doi: 10.1007/s11269-017-1774-7.

[59] J. Adamowski, H. Fung Chan, S. Prasher, B. Ozga-Zielinski, and A. Sliusarieva, "Comparison of multiple linear and nonlinear regression, autoregressive integrated moving average, artificial neural network, and wavelet artificial neural network methods for urban water demand forecasting in Montreal, Canada," ed: Water Resources Research, 48(1), n/a-n/a., 2012.

[60] S. Alvisi, M. Franchini, and A. Marinelli, "A short-term pattern-based water demand forecasting model," *Journal of Hydroinformatics*, vol. 9, pp. 39-50, 01/01 2007, doi: 10.2166/hydro.2006.016.

[61] J. Caiado, "Performance of Combined Double Seasonal Univariate Time Series Models for Forecasting Water Demand," *Journal of Hydrologic Engineering - J HYDROL ENG*, vol. 15, 03/01 2010, doi: 10.1061/(ASCE)HE.1943-5584.0000182.

[62] A. Lomet, F. Suard, and D. Chèze, "Statistical Modeling for Real Domestic Hot Water Consumption Forecasting," *Energy Procedia*, vol. 70, 05/31 2015, doi: 10.1016/j.egypro.2015.02.138.

[63] M. Bakker, J. H. G. Vreeburg, K. M. van Schagen, and L. C. Rietveld, "A fully adaptive forecasting model for short-term drinking water demand," *Environmental Modelling & Software*, vol. 48, pp. 141-151, 2013/10/01/ 2013, doi: https://doi.org/10.1016/j.envsoft.2013.06.012.

[64] E. Pacchin, S. Alvisi, and M. Franchini, "A Short-Term Water Demand Forecasting Model Using a Moving Window on Previously Observed Data," *Water*, vol. 9, no. 3, doi: 10.3390/w9030172.

[65] C. J. Hutton and Z. Kapelan, "A probabilistic methodology for quantifying, diagnosing and reducing model structural and predictive errors in short term water demand forecasting," *Environmental Modelling & Software*, vol. 66, pp. 87-97, 2015/04/01/ 2015, doi: https://doi.org/10.1016/j.envsoft.2014.12.021.

[66] F. Gagliardi, S. Alvisi, Z. Kapelan, and M. Franchini, "A Probabilistic Short-Term Water Demand Forecasting Model Based on the Markov Chain," *Water*, vol. 9, no. 7, doi: 10.3390/w9070507.

[67] C. Copeland, G. MacKerron, and T. J. Foxon, "Futures for Findhorn: Exploring challenges for achieving net zero in an ecological intentional community," *Futures*, vol. 149, p. 103155, 2023/05/01/ 2023, doi: https://doi.org/10.1016/j.futures.2023.103155.

[68] R. J. Hyndman and G. Athanasopoulos, *Forecasting Principles and Practice*. Monash University, Australia, 2018.

[69] G. Shmueli and K. C. J. Lichtendahl, *Practical Time Series Forecasting with R: A Hands-On Guide*, 2 ed. Maryland, USA.: Axelrod Schnall Publishers, 2016.

# 8.0 Appendices

## 8.1 Appendix A: Simple Moving Average Script

```
import pandas as pd
import numpy as np
import matplotlib.pyplot as plt
from sklearn.metrics import mean_squared_error, mean_absolute_error
from math import sqrt

# Load the data
df = pd.read_excel('Summer_hourly_energy_data.xlsx')

# Create a single dataset and calculate SMA on it
df_all = df[['Hour', 'Flat 1 energy (kWh)', 'Flat 2 energy
(kWh)']].copy()
df_all.columns = ['Hour', 'Flat 1', 'Flat 2']
df_all['SMA_5'] = df_all['Flat 1'].rolling(window=5).mean()

# Split into training and test sets again
df_train = df_all[['Hour', 'Flat 1', 'SMA_5']].copy()
df_test = df_all[['Hour', 'Flat 2', 'SMA_5']].copy()

# Plotting the results
fig, ax = plt.subplots(2, 1, figsize=(12,10))

# Training data
df_train[['Flat 1', 'SMA_5']].plot(kind='bar', ax=ax[0])
ax[0].set_title('Training Data')
ax[0].set_xlabel('Hour')
ax[0].set_ylabel('Energy Consumption')

# Test data
df_test[['Flat 2', 'SMA_5']].plot(kind='bar', ax=ax[1])
ax[1].set_title('Test Data')
ax[1].set_xlabel('Hour')
ax[1].set_ylabel('Energy Consumption')

plt.tight_layout()
plt.show()

# Compute the RMSE and MAE for training data (only for those hours
where SMA_5 is not NaN)
rmse_train = sqrt(mean_squared_error(df_train['Flat 1'].dropna()[4:],
df_train['SMA_5'].dropna()))
mae_train = mean_absolute_error(df_train['Flat 1'].dropna()[4:],
df_train['SMA_5'].dropna())

# Compute the RMSE and MAE for test data
```

```
rmse_test = sqrt(mean_squared_error(df_test['Flat 2'][4:],
df_test['SMA_5'][4:]))  # Only compute after the 4th hour
mae_test = mean_absolute_error(df_test['Flat 2'][4:],
df_test['SMA_5'][4:])       # Only compute after the 4th hour

print(f'Training Data RMSE: {rmse_train}')
print(f'Training Data MAE: {mae_train}')
print(f'Test Data RMSE: {rmse_test}')
print(f'Test Data MAE: {mae_test}')
```

## 8.2 Appendix B: AIC and BIC Script

```
import pandas as pd
from statsmodels.tsa.arima.model import ARIMA
import warnings

warnings.filterwarnings('ignore') # Ignore warning messages for clarity

# Load your data
df = pd.read_excel('Summer_hourly_energy_data.xlsx')
data = df['Flat 1 energy (kWh)'].dropna()

# Define the p, d and q values to take on
p_values = [0, 1, 2, 3, 4]
d_values = [0, 1]  # Given data is already differenced once, only
consider 0 or 1 here
q_values = [0, 1, 2, 3, 4]

best_aic = float('inf')  # Start with infinity as initial best
best_bic = float('inf')  # Similarly for BIC
best_order = None

for p in p_values:
    for d in d_values:
        for q in q_values:
            order = (p, d, q)
            try:
                model = ARIMA(data, order=order)
                results = model.fit()
                if results.aic < best_aic:
                    best_aic = results.aic
                    best_bic = results.bic
                    best_order = order
            except:
                continue

print(f'Best ARIMA Order: {best_order}')
print(f'Best AIC: {best_aic}')
print(f'Best BIC: {best_bic}')
```

## 8.3 Appendix C: ARIMA Daily Forecasting Script

```
# 1. Import necessary libraries
import pandas as pd
import numpy as np
import matplotlib.pyplot as plt
from statsmodels.tsa.arima.model import ARIMA
from sklearn.metrics import mean_squared_error, mean_absolute_error

# 2. Load dataset
data = pd.read_excel("Summer_hourly_energy_data.xlsx")

# 3. Create walk-forward validation loop
train = data["Flat 1 energy (kWh)"].values
test = data["Flat 2 energy (kWh)"].values

predictions = []
history = list(train)

for t in range(len(test)):
    # 4. Fit ARIMA model and predict
    model = ARIMA(history, order=(1,0,0))
    model_fit = model.fit()
    output = model_fit.forecast()
    yhat = output[0]
    predictions.append(yhat)
    obs = test[t]
    history.append(obs)

# 5. Store actual and predicted values
actual = pd.Series(test)
predicted = pd.Series(predictions)

# 6. Plot bar plot comparing actual and predicted values for Flat 2
bar_width = 0.35
index = np.arange(len(actual))

fig, ax = plt.subplots(figsize=(12, 7))
bar1 = ax.bar(index, actual, bar_width, label="Actual (Flat 2)",
color="blue")
bar2 = ax.bar(index + bar_width, predicted, bar_width,
label="Predicted", color="red")

ax.set_xlabel("Hour")
ax.set_ylabel("Energy (kWh)")
ax.set_title("Actual vs Predicted Energy Consumption for Flat 2")
ax.set_xticks(index + bar_width / 2)
ax.set_xticklabels([str(i) for i in range(24)])
ax.legend()
plt.show()
```

```
# 7. Calculate and print RMSE and MAE
rmse = np.sqrt(mean_squared_error(actual, predicted))
mae = mean_absolute_error(actual, predicted)
print(f"RMSE: {rmse}")
print(f"MAE: {mae}")

# 8. Create a dataframe to store the actual, predicted values, errors
and MAE, RMSE
results = pd.DataFrame({
    'Actual': actual,
    'Predicted': predicted,
    'Error': actual - predicted,
    'Absolute Error': np.abs(actual - predicted)
})

# Add the RMSE and MAE to the bottom of the dataframe
results = pd.concat([results, pd.DataFrame([[np.nan, np.nan, rmse,
mae]], columns=['Actual','Predicted','Error','Absolute Error'])])

# Print the results
print(results)

# Write to an Excel file
results.to_excel('results1.xlsx')
```