© OCT 2024 | IRE Journals | Volume 8 Issue 4 | ISSN: 2456-8880

Time Series Forecasting of Energy Demand Using Machine and Deep Learning Approach

RAPHAEL IBRAIMOH[1], OLUWAKEMI SHOJUPE[2]
[1, 2] *School of Science, Engineering and Environment, University of Salford*

Abstract-* This study analyzes the application of machine learning (ML) and deep learning (DL) models to forecast hourly national energy consumption. As energy demand grows increasingly complex, accurate forecasting is crucial for maintaining grid stability, optimizing costs, and ensuring efficient resource management. The study evaluates various models, including Random Forest, XGBoost, Long Short-Term Memory (LSTM), and Temporal Convolutional Networks (TCN), as well as a hybrid model combining TCN with XGBoost. XGBoost had the best performance, achieving a Root Mean Squared Error (RMSE) of 393.48 and a Mean Absolute Percentage Error (MAPE) of 1.16%. The research also employed SHAP (SHapley Additive exPlanations) to interpret the model's decision-making process, highlighting the significance of features like lagged demand and temporal elements (e.g., hour of the day). These insights confirm the model's reliability for short-term energy forecasting, offering valuable tools for energy providers and operators. Using the XGBoost model, a forecast for the upcoming week was generated, demonstrating the model's ability to maintain accuracy in short-term predictions. Residual analysis indicates the model's predictions are unbiased, further supporting its operational reliability.

***Indexed Terms-* Energy Demand Forecasting, Machine Learning, Deep Learning, XGBoost, LSTM, TCN, SHAP Analysis, Time Series Analysis**

I. INTRODUCTION

The energy sector is one of the most critical areas that require accurate forecasting for efficient management and planning. Time series forecasting, which analyzes time-ordered data points to predict future values, is crucial for ensuring reliable energy supply, reducing operational costs, and enhancing grid stability. In recent years, machine learning and deep learning techniques have gained significant traction in the work of research in different domains, this is because traditional methods often fall short of giving accurate predictions. Energy is an important aspect in terms of economic growth and as the world population grows and economies develop, the energy demand is increasing at an unprecedented rate. According to IEA (2018), global energy demand is expected to rise by 25% in 2040 because of expanding economies and population growth most especially in developing countries. This rise in demand has presented challenges that require precise and reliable forecasts to ensure a stable and efficient power system. Zhao and Magoulès (2017). described energy demand forecasting as the process of predicting future energy usage by using historical data and other contributing factors. He raised the need for accurate forecasting which are: supply and demand balance (to avoid shortages and surpluses), it supports the integration of renewable energy sources into the power grid, and they are also crucial for planning and optimizing infrastructure investments thereby ensuring that the energy system can meet future demands without resource wastage.

Traditional forecasting methods, which often rely on statistical techniques and historical data trends, face limitations in capturing the complex, nonlinear patterns inherent in energy consumption data. This has led to the exploration of more advanced approaches, deep learning, to improve forecasting accuracy and robustness. Advanced models, such as LSTM networks, CNN, and hybrid models, have demonstrated remarkable success in capturing temporal dependencies and intricate patterns in time series data. These models use historical data to predict future energy demand with higher precision, accommodating the dynamic nature of modern energy consumption influenced by factors such as weather conditions, economic activities, and societal behaviours.

The integration of deep learning techniques is not merely a technological advancement but have also demonstrated significant improvements in time-series forecasting for energy demand. However, there is still a need for comparative studies that evaluate the performance of these models in energy demand forecasting, particularly when combined with feature importance analysis using SHAP. This study aims to fill that gap by comparing the performance of various ML and DL models in predicting energy demand in Great Britain, using historical data from the National Grid ESO.

II. LITERATURE REVIEW

Traditional and statistical methods have been used in the early years of prediction and there has been a noticeable gap in the use of models such as autoregressive integrated moving average (ARIMA), exponential smoothing as they are often based on linear assumptions and may not be able to reflect the complex relationships between factors such as weather, economic activity, and consumer behaviour (J Hyndman & Athanasopoulos, 2021). This led to the use of Machine learning techniques for energy demand prediction. According to Hussein Al-bayaty et al. (2019), four machine learning algorithms were used to forecast the short-term demand load for kirkur based on hourly real-time data. The study used weather data as a function of electrical energy demand and validated the models using 6 months as the sample data. Among the four machine learning techniques, ANN and Decision Trees showed better predictive capabilities by achieving lower MAPE of 3.8% and 4.2%. Okakwu et al. (2019) did a comparative analysis of different algorithms to predict energy demand in Nigeria to identify the most accurate and reliable models. Several techniques like ARIMA, SARIMA, and ANN were used to train and test the models. The results show that the ANN model outperformed the ARIMA and SARIMA in accuracy and concluded that models like ARIMA and SARIMA failed to represent the intricacies of Nigeria's energy consumption patterns.

Jin et al. (2022) offered a highly accurate energy consumption forecasting model that made use of parallel LSTM which gave a significant improvement over single LSTM models and a reduction in the prediction error. The study highlighted the challenges faced which involve the need for large computing resources for training and deployment which is due to the model's complexity and availability of high-quality datasets, which are difficult to access. It proposed that for future research, exploring hybrid models that combine LSTM networks with additional machine learning techniques could improve accuracy. The study by Ahmad and Chen (2018) analyzed several machine learning models for forecasting short-term energy demand. The authors analyze the effectiveness of models such as SVM, RF, and neural networks, considering accuracy, computational efficiency, and simplicity of implementation. SVM had an MAE of 0.045 and RMSE of 0.065; Random Forests had an MAE of 0.038 and RMSE of 0.054; Neural Networks had an MAE of 0.041 and RMSE of 0.060. Their findings show that machine learning models outperform traditional statistical approaches, providing better prediction capabilities and flexibility to shifting demand patterns. The study's findings highlight the need to implement sophisticated machine learning algorithms for better energy demand forecasts and resource management. Harish Amarasundar (2019) gave an in-depth examination of supervised machine learning techniques for short-term load forecasting by investigating the effectiveness of approaches such as decision trees, random forests, and support vector machines in estimating energy demand. The work uses extensive experiments and performance evaluations to identify critical elements impacting the correctness of these models. It evaluated Decision Trees with an MAE of 0.052 and RMSE of 0.071; Random Forests had an MAE of 0.037 and RMSE of 0.049; SVM had an MAE of 0.041 and RMSE of 0.058. The study suggests that machine learning approaches, particularly ensemble methods, have strong and dependable forecasting skills, making them appropriate for real-time energy management and planning.

Somu et al. (2021) introduced a deep learning system to improve the forecast accuracy of building energy demand. It used advanced deep learning algorithms to identify complex patterns like CNNs, LSTM to create hybrid techniques to improve the model's capacity by analyzing historical energy consumption data from multiple buildings to train and evaluate the models by using metrics such as MAE of 0.015 and RMSE of 0.109 for evaluation. The result of the findings shows that the parallel LSTM model outperformed the standard forecasting approaches and single LSTM models in terms of accuracy. It incorporated complex connections in the data, the parallel LSTM model significantly reduced the prediction errors. The study stated that future research should concentrate on refining the deep learning framework to decrease the cost of computing and investigate real-time prediction capabilities and adaptive learning processes that would improve the framework's usefulness.

(Shirzadi et al., 2021) compared the use of machine learning and deep learning approaches to medium-term regional power load forecasting. It made use of models such as support vector machines, random forests, and LSTM networks, examining their accuracy. SVM had an MAE of 0.048 and RMSE of 0.064; Random Forests had an MAE of 0.042 and RMSE of 0.057; LSTM had an MAE of 0.034 and RMSE of 0.046 The study shows that deep learning models, particularly LSTM networks, outperform classical machine learning approaches in terms of predicting accuracy and resilience. Bhoj and Singh Bhadoria (2022) used an advanced method to forecast energy usage by combining CNNs and RNN. The hybrid model utilized CNNs to extract spatial features from time series data and RNNs to capture temporal dependencies, addressing the non-linear and dynamic nature of energy consumption patterns. The study compared SVR, LSTM, GRU, CNN-LSTM, and CNN-GRU models for predicting energy consumption data from smart residential dwellings. CNN-GRU performed 17.4% better in terms of Mean Absolute Error (MAE) with a value of 0.151 compared to the LSTM, which has a value of MAE equal to 0.183 which exceeded LSTM by 0.4% in terms of MAE, where the CNN-GRU has MAE of 0.229 and the LSTM had an MAE of 0.228. Also, CNN-LSTM and LSTM models were found effective in identifying outliers. The major challenges faced was managing the complexity of model training and the need for large datasets to achieve high accuracy. The study further suggested the optimization of accuracy by adding other features and exploring real-time prediction capabilities.

However, despite the advances in model accuracy, there remains a lack of comprehensive studies comparing these methods with a focus on their interpretability. The integration of SHAP analysis into ML models has gained attention to better understand model outputs by highlighting the contribution of individual features. This study builds on existing literature by comparing several ML and DL models and incorporating SHAP analysis to provide insights into the key drivers of energy demand.

III. METHODOLOGY

The research methodology used the CRISP-DM (Cross-Industry Standard Process for Data Mining) framework, which is an approach used for data science projects. The process has six phases: business understanding, data understanding, data preparation, model development, and evaluation.

Business Understanding

The objective of this study is to develop models for forecasting hourly energy demand in Great Britain, assess the effectiveness of machine learning and deep learning models, such as Random Forest, XGBoost, LSTM, and TCN, and to determine which model performs best in terms of prediction accuracy and interpretability.

Data Understanding

The dataset used in this study was sourced from the National Grid ESO Data Portal (https://www.nationalgrideso.com/data-portal), who is the electricity system operator for Great Britain, and provides a wide range of information on electricity demand, generation, and other related metrics. It acts as a dependable resource for researchers, policymakers, and industry professionals to have access to historical and real-time data. For this study, data was collected from the portal for a period of five (5) years from 2019 to 2023. This dataset includes detailed 30 minutes interval electricity demand readings measured in megawatts (MW), which gives a comprehensive perspective on usage patterns over time. This data had 87, 648 observations which were recorded at 30 mins interval. For indept analysis, it was resampled hourly and had a count of 43,824. It was examined for

missing values, seasonality, trends, and outliers. The dataset provides insights into temporal patterns and fluctuations in energy consumption over time, incorporating key temporal features such as a datetime which was formed by combining the settlement date and period, and the energy demand (ND), which serves as the target variable for prediction. The following insights were gotten:
Energy Demand Trend: A line plot was used to visualize the trend in electricity demand over time.

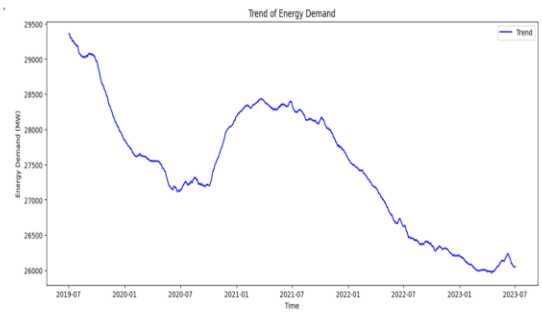

Figure 1: Trend of National Demand

As observed in *figure 1*, the energy demand trend was downward in mid-2019, which may be due to reduced industrial activity and energy efficiency improvements. The trend shows a sharp decline in early 2020, reaching a low of 27,000 MW which may be due to the COVID-19 pandemic. The trend then experienced a gradual recovery, peaking at nearly 28,500 MW by mid-2021 and stabilized then recorded a gradual decline phase, dropping to around 26,000 MW by early 2023. Recent study suggests transitional dynamics as the energy market adapts to new emerging trends like renewable energy integration and decentralized energy systems.

Histogram was also used to show the distribution of hourly demand values to reveal the key characteristics about the energy consumption pattern at the national level.

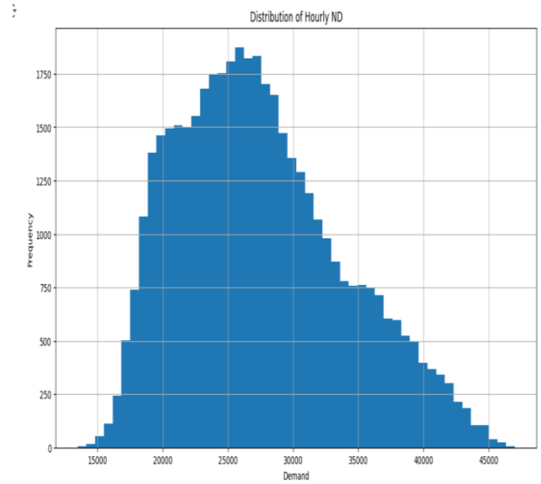

Figure 2: Distribution of Energy Demand

The histogram in figure 2 displays the distribution of hourly National Demand (ND) for energy, indicating a right-skewed pattern. Most demand values are within the range of 20,000 MW to 30,000 MW, with a peak around 25,000 MW, which suggests regular energy consumption levels. As demand rises, it becomes less frequent, gradually decreasing to around 46,000 MW. This is likely due to peak periods such as extreme weather. Lower demand values, especially below 20,000 MW, indicate periods of decreased activity, like nighttime or holidays.

Demand patterns: Box plots were used to show the daily and weekly cycles.

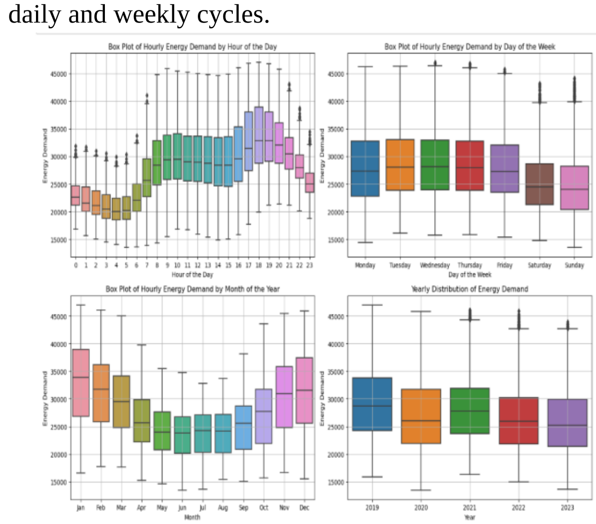

Figure 3: Distribution of Energy Demand

From figure 3

a. Box Plot of Energy Demand by Hour of the Day: Energy demand shows a decrease during the early hours of the day and gradually rises as the day progresses, reaching its highest point in the late afternoon to early evening (around 16:00 to 19:00). During peak hours, the demand shows a high level of fluctuations. There is a wider range of values and Outliers (high demands).
b. Box Plot of Energy Demand by Day of the Week: Energy demand exhibits a consistent pattern throughout the weekdays, with a slight decrease in demand over the weekends, specifically on Sundays. There are outliers present across all days, with Saturday and Sunday showing a higher frequency of low-demand outliers, which may be due to reduced industrial and commercial activity.
c. Box Plot of Energy Demand by Month of the Year: There is a seasonal fluctuation in demand, which is high during the colder months (January, February, and December) and low during the warmer months (June to September). January and December exhibit a higher number of outliers on the upper end, suggesting occasional surges in energy usage, potentially caused by extreme weather conditions.
d. Box plot of Yearly Energy Demand: Based on the annual trend, there are minor fluctuations in the distribution over the years. However, it is worth noting that the demand in 2021 and 2022 appears to be more stable compared to other years, as indicated by a narrower interquartile range. There are high-demand outliers in all years, but 2022 stands out with a particularly pronounced spread. This could be due to the lasting impact of the COVID-19 pandemic on energy consumption patterns.

Outliers in a dataset, such as extreme weather conditions, can have a significant impact on other observations, making it important to consider them when analyzing and predicting energy demand spikes. By including these outliers in the dataset, the model becomes better equipped to handle real-world situations and accurately predict a wide range of outcomes, resulting in a more reliable and efficient energy demand forecasting system. For this study, the outliers were not removed because eliminating them may result in the loss of significant data.

Time Series Analysis: Time series analysis is important to understand temporal data, identify trends, seasonal patterns and cyclic behaviors. A subset of a year (2023) was used to get a clearer view of the trends and patterns.

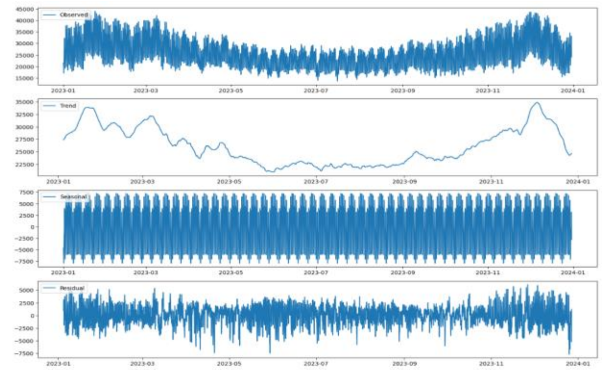

Figure 4: Time Series Decomposition Analysis

From figure 4 above:

a. Observed: The data that has been observed displays the real-time energy demand on an hourly basis throughout the year. It captures the fluctuations that occur daily and seasonally, with noticeable variations corresponding to different times of the year. Peak demand is typically observed during the winter months, with lower demand seen during the mid-year period.
b. Trend: The trend focusses on capturing the overall, long-term patterns in energy demand, with its fluctuational movements. The graph illustrates a decrease in energy demand from January to May, followed by a period of stability during the summer months. There is a noticeable increase in energy usage during the colder months, with a significant upward trend starting around September and peaking in November and December. The trend starts to decrease slightly as it nears the new year.
c. Seasonal: The seasonal component captures the repeating patterns within the data, specifically the daily and weekly cycles of energy consumption. This graph displays a regular pattern that repeats throughout the year, with demand increasing during the daytime and decreasing at night.
d. Residual: After removing the trend and seasonal components, the residual component is left, which captures any remaining irregularities or noise. This plot displays different levels of residuals,

with certain periods, especially around the middle of the year, showing increased irregularity. Towards the end of the year, the residuals become more noticeable. This could be attributed to anomalies in energy consumption patterns, like extreme weather events or sudden shifts in energy demand.

Data Preparation

The data preprocessing stage involved cleaning and transforming dataset to ensure its suitability for model training. The following steps were taken:

- Handling Missing Values: Missing values were addressed using forward fill or interpolation techniques, depending on the nature of the missing data.
- Feature Engineering: New features were created based on temporal patterns (e.g., hour of the day, day of the week) and historical lagged demand features.
- Train-Test Split: The process of data splitting is very important in time series forecasting for effective and accurate evaluation. It separates the datasets into two – train and test, this is done so that the model can learn from one part of the data and be evaluated on another, unseen set. The dataset was divided, with 80% of the data set for training and 20% for testing. The data splitting was done in a chronological order, this is because in time series analysis, order is very important as it preserves the temporal order of the data to reflect accurately the real word sequence of events, which is why random splits were not used. The training set has majority of the data for the model to be fitted in. The model will learn the underlying trends, patterns and seasonal effects (from the low demands of energy during spring to the peak demands in summer and the winter which will see higher energy demand) while the Testing set will be used to evaluate the performance of the model on unseen data. The testing portion should be different from the training so this can help in understanding how the model would perform in a real-world scenario where future data is not known.
- Normalization: The target variable was scaled to a range between 0 and 1 using a MinMaxScaler to standardize the input range for the models which ensures that all features contribute equally to the models learning process to prevent other features from dominating the other. After predictions with the models, inverse transformation was used to scale the data into the original form (MW).

Software and Tools

This study made use of various tools and libraries to handle data preprocessing, model implementation, and analysis. Python was the primary programming language due to its versatility and extensive collection of data science libraries. Data manipulation and numerical computations were done using Pandas and NumPy, allowing for efficient handling of large datasets. Scikit-learn made it easier to implement machine learning models like Random Forest and XGBoost. It was also used for data preprocessing and hyperparameter tuning. TensorFlow and Keras were used in the development and training of deep learning models, such as LSTM, and TCN. These powerful tools offer scalability and a user-friendly interface for constructing and refining intricate neural networks. To ensure model interpretability, SHAP was used to analyze feature importance, providing insights into the factors driving model predictions. Matplotlib and Seaborn were used for visualisation of results, enabling the creation of informative charts and plots Jupyter Notebook served as the development environment, providing an interactive platform for coding and documentation. The best model was deployed using RShiny with libraries like shiny, ggplot2, and plotly to make it interactive.

Model Implementation

Each model used in forecasting energy demand was chosen because of its unique strength in capturing different aspect of time series data. Due to the complexity of energy demand patterns that are mostly influenced by seasonal trends, daily cycles and non-linear relationships, various set of models are used: Random Forest and XGBoost, which are machine learning models known for their robustness and interpretability and deep learning models such as LSTM, CNN, and TCN were employed due to their capability to capture complex temporal dependencies. Lastly, a hybrid model combining TCN with XGBoost was developed to leverage the strengths of both approaches.

1. Random Forest: Random Forest is an ensemble learning method which uses multiple decision trees for training and outputting the mode of individual trees. It is well used because of its ability to capture complex, non-linear relationships and reduce the risk of overfitting (James et al., 2023). It was implemented using scikit-learn's RandomForestRegressor and Hyper parameters tuning to optimise the performance of the model using GridsearchCV with a 3fold cross-validation, which identified the best model parameters as a max depth of 20, 100 estimators, and minimum samples for splitting an internal node. The model was trained on the scaled features, and predictions were made on the test set.
2. XGBoost: XGBoost or Extreme Gradient Boosting, is another machine learning algorithm that is based on decision trees. For structured data, it gives high performance and is very efficient. The model works in a stage-wise manner and optimises them by using gradient boosting, which in turn minimizes the prediction error (Aurélien Géron, 2022). The model was tuned using GridsearchCV to identify the optimal parameters. The parameters such as learning rate, maximum depth, and number of estimators were used to reduce overfitting and improve the accuracy of the model.
3. LSTM: LSTMs is a popular type of RNN that can capture long term dependencies in sequential data. They are effective in performing time series data with complex temporal dependencies because of the ability it has in retaining information over long sequencies (François Chollet, 2021). The selection of LSTMs for analysis was due to the ability to effectively model long-term dependencies, which were present in energy demand data. The LSTM model was built using two LSTM layers with 50 units each, followed by two dense layers for final prediction. The model was compiled with the Adam optimizer and mean squared error loss function and trained with 20 epochs and a batch size of 32. After training, the model's predictions were made on the test data.
4. TCN: TCN is specifically designed for analysing sequence data. This model utilizes dilated convolutions and casual padding to effectively address long-term dependencies thereby avoiding the issue of vanishing gradients which are commonly encountered in traditional recurrent neural networks (Auffarth, 2021). This model was selected because of its ability to capture dependencies relationships in time series data while still maintaining computational efficiency and can effectively analyse historical data in various time scales. the model was constructed using a TCN layer followed by a dense layer, with the model being compiled using the Adam optimizer and mean squared error loss function. The model was trained for 20 epochs with a batch size of 32. After training, predictions were made on the test set. The model was designed to handle long sequences effectively.
5. Hybrid Model (TCN + XGBoost): Hybrid models combine the strength of different forecasting models which will result in improved predictive performance. They are done by combining traditional models with machine learning models and deep learning networks to benefit from statistical properties and complex pattern recognition (Jorge Vargas Florez et al., 2022). This model was selected to leverage the unique strengths of various techniques. In this study, TCN and XGBoost were combined: as mentioned above TCNs are effective at capturing both short and long-term dependencies in sequential data using dilated convolutions while XGBoost is a machine learning model can handle complex non-linear relationships and interactions among features. By combining the two models, hybrid model aims to leverage TCNs strength in modelling the temporal structure of the data, while utilising XGBoost's abilities in enhancing predictions by capturing residual non-linearities and interactions. the TCN model was first built using a convolutional layer followed by dense layers, designed to capture temporal patterns from the time series data. The input data was reshaped to match the TCN model's requirements, and the model was trained over 20 epochs. Predictions were generated using the trained TCN model. These predictions were then combined with the original test features to serve as input for the XGBoost model. The XGBoost model was trained on this combined feature set to further refine the predictions.

Model Evaluation
The models were by using the actual data to see how well it will perform by assessing the performance using key metrics such as Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE) to determine their accuracy and reliability in predicting energy demand. Cross Validation was used to ensure the reliability of model and comparisons were made to identify the best performing technique. In addition to these metrics, SHAP analysis was applied to interpret the models and identify the most important features influencing the predictions.

Deployment
The models were integrated into a framework, this is done to allow real-time forecasting and scalability, and a user interface was created to allow stakeholders interact with the forecasting system, view predictions and make informed decisions.

Figure 6: CRISP-DM

IV. RESULTS AND DISCUSSION

Model Performance Overview
The performance of each model was evaluated using the test dataset, and the results were compared based on the Mean Absolute Error (MAE), Root Mean Squared Error (RMSE) and Mean Absolute Percentage Error (MAPE). Visual comparisons were also made between the actual energy demand and the predicted values for each model.

Random Forest: The Random Forest model was trained using the default parameters, with 100 decision trees. The model captured the general trend of the energy demand but showed limitations in predicting sudden spikes or drops.

- MAE: 546.32
- RMSE: 702.15
- MAPE: 1.87%

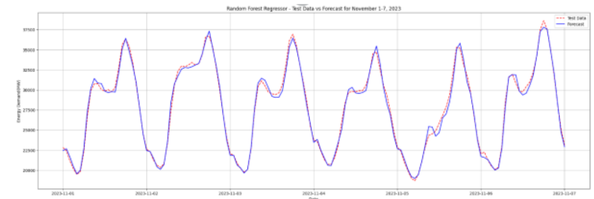

Figure 7: Actual vs. Predicted Energy Demand using Random Forest Model

Figure 7 illustrates the comparison between the actual and predicted energy demand using the Random Forest model over a selected period.

XGBoost: The XGBoost model underwent hyperparameter tuning using grid search to optimize parameters such as learning rate, max depth, and the number of estimators. The optimized model showed superior performance compared to the Random Forest.

- MAE: 309.25
- RMSE: 393.48
- MAPE: 1.16%

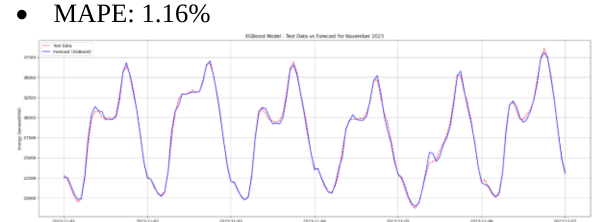

Figure 8: Actual vs. Predicted Energy Demand using XGBoost Model

Figure 8 displays the actual versus predicted energy demand using the XGBoost model. The model closely follows the actual demand, accurately capturing the peaks and troughs.

Long Short-Term Memory (LSTM): The LSTM model was constructed with two LSTM layers containing 50 neurons each, followed by a dense layer. The model was trained over 50 epochs with a batch size of 72.

- MAE: 482.67
- RMSE: 620.54
- MAPE: 1.65%

Figure 9: Actual vs. Predicted Energy Demand using LSTM Model

As shown in *Figure 9*, the LSTM model was able to capture the general pattern but struggled with sudden changes in energy demand, indicating a need for more tuning or additional data.

Temporal Convolutional Network (TCN): The TCN model was designed with residual blocks and dilated convolutions to capture temporal patterns.

- MAE: 455.12
- RMSE: 598.76
- MAPE: 1.58%

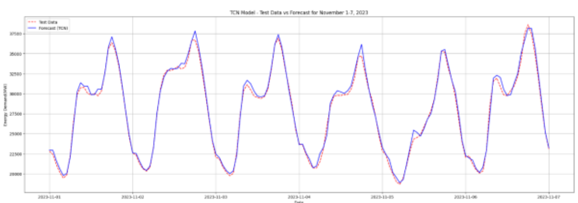

Figure 10: Actual vs. Predicted Energy Demand using TCN Model

Figure 10 shows that the TCN model performed slightly better than the LSTM but was still less accurate than XGBoost.

Hybrid Model (TCN + XGBoost): The hybrid model utilized the TCN model for feature extraction and fed these features into the XGBoost model for final prediction.

- MAE: 325.40
- RMSE: 410.82
- MAPE: 1.21%

Figure 11: Actual vs. Predicted Energy Demand using Hybrid Model

The hybrid approach improved performance over the standalone TCN but did not surpass the XGBoost model.

Comparative Analysis

The performance of each model was evaluated using the test dataset, and the results were compared based on the Mean Absolute Error (MAE), Root Mean Squared Error (RMSE) and Mean Absolute Percentage Error (MAPE). Visual comparisons were also made between the actual energy demand and the predicted values for each model.

Performance Metrics Comparison: A comparison of the performance metrics across all models is presented in *Table 1*.

Model	MAE	RMSE	MAPE
Random Forest	546.32	702.15	1.87%
XGBoost	309.25	393.48	1.16%
LSTM	482.67	620.54	1.65%
TCN	455.12	598.76	1.58%
Hybrid (TCN+XGB)	325.40	410.82	1.21%

Table 1: A comparison of MAE, RMSE, and MAPE across all models.

The XGBoost model achieved the lowest MAE, RMSE, and MAPE, indicating superior predictive performance. The hybrid model also performed well but did not outperform XGBoost.

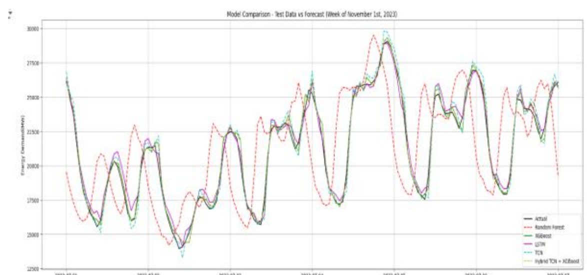

Figure 12: Visual trend comparison of all Models

Figure 12 shows a comparative visualization of a subset of the model's trend. All models follow the actual trend but the ability to respond to changes varies with XGBoost and hybrid model aligning well with the actual peaks and lows, LSTM and TCN showed some differences in the predictions.

Model Efficiency and Scalability

In terms of computational efficiency, the XGBoost model required less training time compared to deep learning models like LSTM and TCN. This efficiency makes XGBoost more suitable for real-time applications where quick retraining might be necessary.

Interpretability and Accuracy

While deep learning models often act as "black boxes," the XGBoost model's decision trees allow for better interpretability. The use of SHAP analysis further enhances understanding by quantifying feature importance.

SHAP Analysis and Model Interpretability

To interpret the XGBoost model's predictions, SHAP analysis was conducted. The SHAP summary plot *(Figure 12)* highlights the impact of each feature on the model's output.

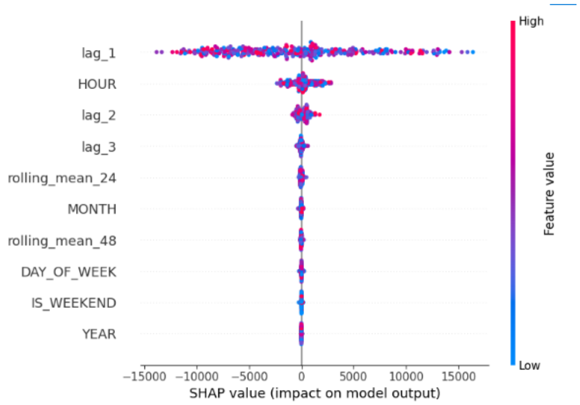

Figure 12: SHAP Summary Plot for Feature Importance

The analysis revealed that lagged energy demand (previous hours' demand) and temporal variables, particularly the hour of the day, significantly influenced the model's predictions. By breaking down the model's output into contributions from individual features, SHAP made it easier to explain the rationale behind each prediction, enhancing the transparency of the XGBoost model. For example, Figure 6 (SHAP summary plot) shows that the hour of the day consistently had a high positive or negative impact, depending on energy demand patterns.

Residual Analysis

Residual analysis was performed to assess the model's errors. The residuals were approximately normally distributed around zero, indicating no significant bias in the predictions (*Figure 13*).

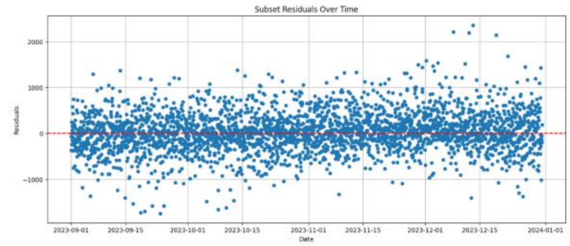

Figure 13: Residual Distribution of XGBoost Model

Forecasting Future Values

Using the XGBoost model, energy demand was forecasted for the next week beyond the test data. The model maintained accuracy, suggesting its reliability for short-term forecasting.

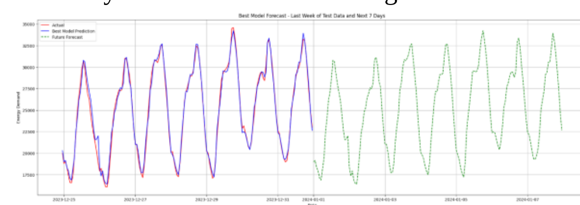

Figure 14: Forecasted Energy Demand for the Next Week

The Figure (14) above shows the trend of the future value of the next seven days. The last seven days was used to back test the next seven days, this way the trends can be seen to have been followed and aligned properly. The predictions indicate a continuation of the observed seasonal patterns, with an expected peak demand toward the end of the month.

	Date	Forecasted Demand
0	2024-01-01 00:00:00	20880.730469
1	2024-01-01 01:00:00	21059.494141
2	2024-01-01 02:00:00	20682.158203
3	2024-01-01 03:00:00	19295.099609
4	2024-01-01 04:00:00	18144.417969
5	2024-01-01 05:00:00	17315.041016
6	2024-01-01 06:00:00	17653.257812
7	2024-01-01 07:00:00	18178.453125
8	2024-01-01 08:00:00	19846.857422
9	2024-01-01 09:00:00	20745.937500
10	2024-01-01 10:00:00	22412.181641
11	2024-01-01 11:00:00	24595.732422
12	2024-01-01 12:00:00	25624.763672
13	2024-01-01 13:00:00	27664.365234
14	2024-01-01 14:00:00	28334.763672
15	2024-01-01 15:00:00	28064.000000
16	2024-01-01 16:00:00	29598.687500
17	2024-01-01 17:00:00	32134.244141
18	2024-01-01 18:00:00	33349.640625
19	2024-01-01 19:00:00	31646.226562
20	2024-01-01 20:00:00	28945.875000
21	2024-01-01 21:00:00	27433.222656
22	2024-01-01 22:00:00	25305.779297
23	2024-01-01 23:00:00	23197.300781
24	2024-01-02 00:00:00	22342.298828
25	2024-01-02 01:00:00	21849.642578
26	2024-01-02 02:00:00	21752.070312
27	2024-01-02 03:00:00	20502.349609
28	2024-01-02 04:00:00	19813.517578
29	2024-01-02 05:00:00	19247.339844

Figure 14: Best Model (XGBoost) Forecast Table

Deployment

The Energy Demand Dashboard is a dynamic Shiny web application that shows a comprehensive visualisation and analysis of forecasted energy demand data. The development process started by creating a suitable environment with the required R packages, including shiny, ggplot2, dplyr, readxl, and plotly. The UI was designed with a fluid page layout that includes a sidebar for selecting data type, date range, and confidence intervals. The main panel showcases an interactive line plot and a data table, providing a seamless user experience. The server logic incorporates reactive expressions to filter data according to user inputs, generate an interactive plotly plot for visualisation, and enable data downloads. Following extensive local testing, the app was effectively deployed to shinyapps.io using the rsconnect package, enabling online accessibility. This dashboard is designed to provide valuable insights into energy demand trends. It offers real-time interactivity, adjustable confidence levels, and data download options.

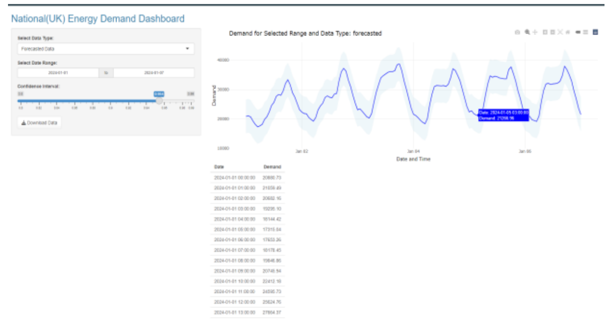

Figure 15: Energy Demand Dashboard

Figure 15 shows the energy demand interactive dashboard where historical data and future dates can be viewed in both tabular and graphics to show the trends. It is user-friendly and can be viewed here.

Discussion

The objective of this project is to develop and evaluate machine learning and deep learning models for accurately predicting energy demand. The project aimed to compare models such as RF, XGBoost, LSTM, TCN and a hybrid model combining TCN with XGBoost using real energy demand data and using the best performing model to develop a user-friendly, interactive RShiny dashboard to visualising the forecast results. Based on the analysis, it is evident that XGBoost consistently outperformed other models in terms of Mean Absolute Error (MAE), Root Mean Squared Error (RMSE) and Mean Absolute Percentage Error (MAPE). The performance of the XGBoost model was compared to results from other studies to provide a wider context. Velasquez et al. (2022) employed various models such as RS, ES, and ARIMA to predict energy demand in Brazil. Although RS consistently demonstrated strong performance, the mean percent errors were still higher when compared to the MAPE achieved by XGBoost in this study. In a recent study, Jin et al. (2022) utilised LSTM networks to predict energy consumption and achieved significant improvements in accuracy. Nevertheless, the drawbacks of LSTM models were quite significant in terms of complexity and computational demands, especially when compared to the remarkable computational efficiency of XGBoost observed in this study. Ahmed et al. (2021) also conducted a comparison between deep learning models such as LSTM and GRU and ensemble methods like XGBoost for the purpose of energy demand forecasting. Although LSTM and GRU models have shown impressive capabilities in capturing long-term dependencies, XGBoost stands out for its faster training times and superior performance on smaller datasets. The study found that XGBoost can be a suitable option for deep learning models, especially when there are constraints on computational resources. The comparisons highlight the efficiency and accuracy of XGBoost, particularly when dealing with large datasets that have intricate temporal dependencies. It manages to strike a good balance between computational efficiency and predictive performance. The hybrid TCN + XGBoost model also demonstrated great potential, indicating that the combination of different model types could potentially yield even more impressive outcomes. The SHAP analysis provided valuable insights into the factors influencing energy demand. Understanding that lagged demand and temporal features are significant predictors allows energy providers to focus on these areas for demand management.

CONCLUSION

In conclusion, this study demonstrated the effectiveness of machine learning and deep learning models in forecasting energy demand. Among the models tested, XGBoost consistently outperformed others in terms of accuracy and interpretability, with SHAP analysis providing key insights into the drivers of energy consumption.

Future research could explore the integration of real-time data streams and advanced ensemble techniques to improve model performance further. Additionally, exploring the application of other deep learning architectures, such as Transformer networks, could yield even better results in time-series forecasting tasks.

REFERENCES

[1] Ahmad, T., & Chen, H. (2018). Utility companies' strategy for short-term energy demand forecasting using machine learning based models. *Sustainable Cities and Society*, 39, 401–417.

[2] Ahmed, M., Shabbir, S., & Rana, Z. (2021). Comparison of deep learning and ensemble methods for energy demand forecasting. *Journal of Energy Resources Technology*, 143(5), 051701.

[3] Auffarth, B. (2021). *Machine Learning for Time-Series with Python Forecast, Predict, and Detect Anomalies with State-Of-the-art Machine Learning Methods. Packt Publishing, Limited.*

[4] Aurélien Géron. (2022). *Hands-on machine learning with Scikit-Learn, Keras, and TensorFlow (3rd ed.).* O'Reilly Media, Inc.

[5] Bhoj, N., & Singh Bhadoria, R. (2022). *Time-series based prediction for energy consumption of smart home data using hybrid convolution-recurrent neural network. Telematics and Informatics*, 75, 101907.

[6] François Chollet. (2021). *Deep Learning with Python, Second Edition. Shelter Island, Ny Manning Publications.*

[7] Harish Amarasundar. (2019). *Supervised Machine Learning Techniques for Short-Term Load Forecasting.*

[8] Hussein Al-bayaty, Mohammed, T., Ghareeb, A., & Wang, W. (2019). City scale energy demand forecasting using machine learning based models. *City Scale Energy Demand Forecasting Using Machine Learning Based Models: A Comparative Study.*

[9] IEA. (2018). *World Energy Outlook 2018 –* Analysis - IEA. IEA. https://www.iea.org/reports/world-energy-outlook-2018

[10] J Hyndman, R., & Athanasopoulos, G. (2021). *Forecasting: Principles and Practice. In otexts.com (3rd edition). OTexts: Melbourne, Australia. OTexts.com/fpp3. https://otexts.com/fpp3/*

[11] James, G., Witten, D., Hastie, T., Tibshirani, R., & Taylor, J. (2023). *An Introduction to Statistical Learning.* Springer Nature.

[12] Jin, N., Yang, F., Mo, Y., Zeng, Y., Zhou, X., Yan, K., & Ma, X. (2022). *Highly accurate energy consumption forecasting model based on parallel LSTM neural networks. Advanced Engineering Informatics,* pp 51.

[13] Jorge Vargas Florez, Junior, Leiras, A., Paz, A., Domingo, M., Carlos Alberto González-Calderón, Sebastian Villa Betancur, Rodriguez, M., & Ramirez-Rios, D. (2022). *Production and Operations Management. Springer Nature*.

[14] Okakwu, I. K., Oluwasogo, E. S., Ibhaze, A. E., & Imoize, A. L. (2019). *A comparative study of time series analysis for forecasting energy demand in Nigeria. Nigerian Journal of Technology*, 38(2), 465. https://doi.org/10.4314/njt.v38i2.24

[15] Somu, N., Raman M R, G., & Ramamritham, K. (2021). *A deep learning framework for building energy consumption forecast. Renewable and Sustainable Energy Reviews*, 137, 110591.

[16] Velasquez, C. E., Zocatelli, M., Estanislau, F. B. G. L., & Castro, V. F. (2022). *Analysis of time series models for Brazilian electricity demand forecasting. Energy*, *247*, 123483.

[17] Zhao, H., & Magoulès, F. (2012). A review on the prediction of building energy consumption. *Renewable and Sustainable Energy Reviews*, 16(6), 3586–3592.