International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 09 Issue: 04 | April - 2025 SJIF Rating: 8.586 ISSN: 2582-3930

Energy Consumption Prediction for Smart Homes

Sourav Sanjay Mane, Paras Atal

Sourav Sanjay Mane, Department of Computational Intelligence, SRMIST
Paras Atal, Department of Computational Intelligence, SRMIST

---***---

Abstract - Accurate electricity consumption forecasting is a critical component in ensuring the stability, efficiency, and cost- effectiveness of national power systems. In Finland, where the electricity sector is shaped by nuclear energy, cogeneration, renewable sources like black liquor and wood, and significant electricity imports, reliable forecasting becomes even more essential. This research investigates the application of Long Short- Term Memory (LSTM), a type of recurrent neural network (RNN) designed for sequential data, in predicting future energy consumption trends. A univariate time series dataset from Fin- grid, Finland's transmission system operator, encompassing six years of hourly electricity consumption data, serves as the basis for model training and evaluation. The study aims to assess the model's performance in capturing temporal patterns and generating accurate short-term forecasts. Our results highlight the capability of LSTM models to effectively model complex energy usage patterns and provide meaningful insights for energy planning and decision-making processes.

***Key Words*:** Long short-term memory (LSTM), Time series analysis, Machine learning, Energy prediction, Fingrid, Univariate data, Finland electricity sector.

1.INTRODUCTION

The electricity sector in Finland is characterized by a diverse energy mix, including nuclear power, the forest industry, black liquor and wood-based fuels, combined heat and power (CHP) production, and electricity imports from neighboring countries (Electricity Sector in Finland, 2022). In 2020, Finland's total electricity production reached 66.6 terawatt-hours (TWh), of which 34.7 TWh—approximately 52%—was generated from renewable energy sources (Statistics, 2021).

Accurate forecasting of electricity consumption plays a critical role in the efficient operation, planning, and economic optimization of power systems. The reliability of demand forecasts directly influences the operational decisions, grid stability, and long-term investments of utility companies.

In recent years, artificial intelligence (AI) and machine learning (ML) techniques have emerged as powerful tools for solving a wide range of predictive and analytical problems. Among these, Long Short-Term Memory (LSTM) networks—a specialized form of recurrent neural networks (RNNs)—have proven particularly effective in modeling time series data due to their ability to capture temporal dependencies and long-range patterns.

This study explores the feasibility and performance of applying LSTM models to forecast electricity consumption in Finland. The model is trained on a univariate time series dataset provided by Fingrid, Finland's transmission system operator, which comprises six years of hourly electricity consumption data. The primary objective is to evaluate the LSTM model's ability to generate accurate and reliable forecasts that can aid in energy planning and policy-making.

2. Fundamentals

2.1 Machine Learning (ML) and Initial Optimizations

Machine Learning (ML) is the foundation of modern artificial intelligence (AI), where systems are trained to make decisions or predictions based on data. Early machine learning models, such as Linear Regression, Support Vector Machines (SVMs), and Decision Trees, are designed to handle structured, tabular data. These algorithms perform well when the relationship between input and output variables is linear or close to linear. However, ML models face several limitations when dealing with more complex datasets:

- Feature Engineering: ML algorithms require extensive feature engineering to extract relevant patterns from raw data.

Limitations in Handling

- Complex Data: ML struggles with high-dimensional data, such as images or sequential data.

- Scalability: With large datasets, traditional ML models face computational bottlenecks, requiring optimization techniques like regularization and cross-validation to reduce overfitting and improve generalization.

Despite these limitations, ML provided the foundational frameworks that deep learning models would later expand upon, particularly by allowing the development of more sophisticated techniques for handling larger and more complex data.

2.2 Deep Learning (DL) and Feature Hierarchies

Deep Learning (DL) represents a leap forward from traditional ML by utilizing neural networks with multiple layers, also known as deep neural networks (DNNs). DL models, such as Convolutional Neural Networks (CNNs) and Fully Connected Networks (FCNs), are capable of learning from large datasets without explicit feature engineering. These models automatically learn feature hierarchies, which is a major advantage over traditional ML models.

The key benefits of DL include:

- Automatic Feature Learning: DL models can automatically extract meaningful features from raw data, which reduces the need for manual intervention.

- Handling Complex Data Types: DL is capable of processing more complex data, such as images, text, and unstructured data, making it well-suited for applications like image recognition and natural language processing.

However, as the models grow deeper, challenges emerge:

- Vanishing Gradient Problem: In very deep networks, the gradients used in backpropagation can shrink to near-zero values, making training difficult.

- Computational Expense: DL models require vast computational resources and large datasets to train effectively, which can be a limitation in certain environments.

While deep learning significantly improved model performance, these challenges led to the need for specialized architectures to handle sequential and time-series data more effectively.

2.3 Recurrent Neural Networks (RNNs): Addressing Sequential Data

Recurrent Neural Networks (RNNs) were developed as a solution to the problem of modeling sequential data, such as time series, speech, and text. Unlike traditional feed-forward networks, RNNs have feedback connections, allowing information to persist in the model, making them ideal for tasks where the current output depends on previous inputs.

RNNs are effective for:

- Modeling Sequential Data: RNNs can process inputs in sequences and maintain context across timesteps.

- Handling Variable-Length Inputs: Unlike traditional ML models that require fixed-size inputs, RNNs can handle inputs of varying lengths, making them versatile for many real-world applications.

However, RNNs still face some critical challenges:

- Vanishing Gradient Problem: When training on long sequences, the gradients can diminish, making it difficult for the network to learn long-term dependencies.

- Exploding Gradient Problem: In contrast, the gradients can also become too large, causing instability during training.

Despite these issues, RNNs represented a significant advancement in deep learning, especially for applications requiring the processing of temporal or sequential data.

2.4 Long Short-Term Memory (LSTM): Overcoming RNN Limitations

To address the limitations of traditional RNNs, Long Short-Term Memory (LSTM) networks were introduced. LSTMs are a specialized type of RNN designed to capture long-term dependencies and mitigate the vanishing and exploding gradient problems. By introducing gates (input, output, and forget gates), LSTMs are able to control the flow of information more effectively over time, allowing them to remember important information for longer periods.

The key improvements of LSTMs over traditional RNNs include:

- Long-Term Memory: LSTMs can capture long-term dependencies in sequential data, allowing them to remember information for much longer than traditional RNNs.

- Stability in Training: The gating mechanism ensures that gradients remain stable during backpropagation, preventing the vanishing and exploding gradient problems.

- Improved Performance on Sequential Tasks: LSTMs have become the go-to architecture for tasks like speech recognition, machine translation, and time series forecasting, where long-range dependencies are crucial.

LSTMs have revolutionized how we approach sequential data, offering significant improvements over both ML and traditional DL models.

3. Time Series Forecasting

3.1 Time Series

Time series data refers to a sequence of data points measured at successive, evenly spaced points in time. The primary goal of time series analysis is to forecast future values based on previously observed data, identifying patterns and trends over time. Time series can be categorized based on the time variable and dependent variables:

- Continuous Time Series: Data is measured in continuous time, such as temperature measurements.

- Discrete Time Series: Data is measured at specific intervals, such as daily, weekly, or yearly records of energy consumption or population changes.

Time series data can also be split into:

- Univariate Time Series: A single variable is measured over time (e.g., daily stock prices).

- Multivariate Time Series: Multiple variables are measured over time (e.g., forecasting sales based on multiple factors like marketing expenditure and seasonality).

While time series analysis focuses on extracting meaningful statistics from the data, forecasting is aimed at predicting future values by leveraging past observations.

3.2 Time Series Characteristics

Autocorrelation: This measures the relationship between a variable's current value and its past values. High autocorrelation indicates a strong dependency on previous time points.

Seasonality: This represents patterns that repeat at regular intervals (e.g., monthly, quarterly). For example, sales data may exhibit higher values during holiday seasons.

Stationarity: A time series is considered stationary if its statistical properties (mean, variance, and covariance) do not change over time. Non-stationary data needs to be transformed (e.g., differencing) to become stationary for better forecasting accuracy.

3.3 Time Series Forecast Methods

Time series forecasting aims to predict future values by learning from historical patterns in the data. Classical time series forecasting methods include:

- Autoregressive Moving Average (ARMA): A model that combines autoregressive and moving average components for stationary time series.

- Autoregressive Integrated Moving Average (ARIMA): An extension of ARMA, which includes differencing to handle non-stationary time series data.

Seasonal ARIMA (SARIMA): An enhancement of ARIMA to capture seasonality in the data.

In addition to these classical methods, machine learning techniques have emerged as powerful tools for time series forecasting. These techniques can better capture non-linearities and complex patterns in the data:

- Multi-Layer Perceptron (MLP): A type of feed-forward neural network used for regression and forecasting.

- Bayesian Neural Networks (BNN): A probabilistic neural network that incorporates uncertainty into predictions, making it more robust.

- Generalized Regression Neural Networks (GRNN): A type of radial basis network that can model complex relationships in the data.

More recently, Recurrent Neural Networks (RNN) and Long Short-Term Memory (LSTM) networks have become popular for time series forecasting due to their ability to capture long-term dependencies and temporal patterns.

3.4 RNN and LSTM for Time Series Forecasting

Recurrent Neural Networks (RNN) are designed to handle sequential data by maintaining hidden states that capture information from previous time steps. This makes RNNs particularly suitable for time series forecasting, where the sequence of past values is crucial for predicting future ones. However, RNNs suffer from the problem of vanishing gradients, making them less effective at capturing long-term dependencies.

Long Short-Term Memory (LSTM) networks were introduced to overcome the limitations of traditional RNNs. LSTMs are designed to retain information for longer periods, making them ideal for modeling time series data with long-term dependencies. By using memory cells and gates (input, forget, and output gates), LSTMs are able to selectively retain and forget information, allowing them to make more accurate predictions over long sequences.

3.5 Comparison of Classical and Modern Time Series Forecasting Techniques

The evolution of time series forecasting from classical statistical models to modern machine learning methods has significantly improved forecasting accuracy and scalability:

Classical Methods (ARMA, ARIMA, SARIMA): These methods are effective for stationary data and datasets with clear linear patterns. However, they often struggle with non-linear relationships, large datasets, and data with complex dependencies.

Machine Learning Methods (MLP, BNN, GRNN): These methods can capture more complex, non-linear patterns but may require more computational resources, especially for large datasets.

RNN and LSTM: These deep learning models excel in capturing long-term dependencies and non-linearities in time series data. They are particularly effective when dealing with large datasets with complex temporal patterns, but they require substantial computational power and are more challenging to interpret.

4. Visualizations

4.1 Data

The objective of this study is to develop a Long Short-Term Memory (LSTM) model for forecasting future energy consumption. The dataset utilized for this purpose was obtained from Finland’s transmission system operator and comprises 52,965 observations across five variables. Initial data quality assessments revealed no missing or duplicate values. The energy consumption values range between 5,341 MWh and 15,105 MWh, with an average consumption of 9,488.75 MWh.

The dataset records start and end timestamps in both Coordinated Universal Time (UTC) and Helsinki local time (UTC+3). For the purposes of this study, the analysis was centered on the Helsinki local time and corresponding energy consumption data.

Comprehensive exploratory data analysis was conducted to understand the underlying structure and trends within the dataset. Visualizations and descriptive statistics indicated a clear seasonal pattern, with energy consumption peaking during the winter months and declining during the summer. The distribution of consumption values appeared unimodal with a slight right skew, and no significant outliers were detected. The dataset spans a six-year period, from 2016 to 2021.

For pre-processing, the dataset was refined to align with complete weekly periods, starting on Monday, 4 December 2016, and concluding on Sunday, 26 February 2021. This adjustment necessitated the removal of 71 observations at the beginning and 121 observations at the end of the dataset. The final dataset retained two primary columns, 'Date Time' and 'Consumption,' from which additional temporal features such as Month, Year, Date, Week, and Day were extracted to enhance subsequent analysis and visualization processes. Null value analysis confirmed the absence of missing entries, thus validating the dataset’s integrity for use in model development and evaluation.

4.2 Graphs

The hourly analysis of energy consumption reveals distinct intraday patterns. During the early morning hours (00:00–05:00), consumption remains relatively low, reflecting minimal residential and industrial activity. A noticeable increase is observed starting around 06:00, corresponding to the beginning of daily operations in households and workplaces. Throughout the midday hours (10:00–17:00), consumption levels remain relatively stable, coinciding with typical working periods. In the evening (18:00–21:00), there is a secondary rise in energy usage, likely driven by residential demand as people return home, followed by a gradual decline during the late evening and nighttime.

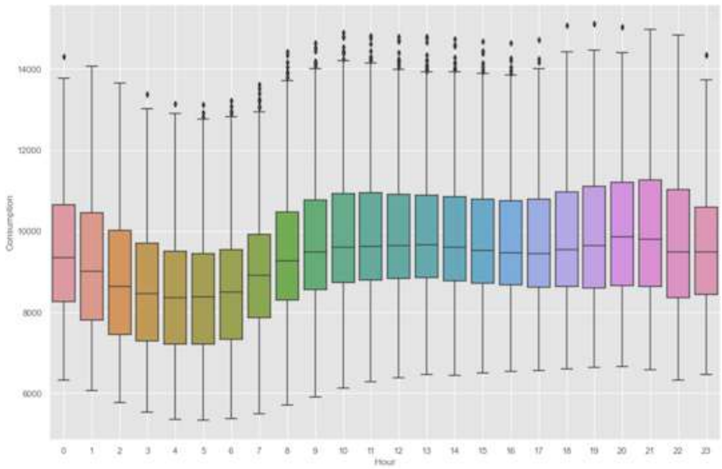

Fig -1: Energy consumption VS Hour

The hourly analysis of energy consumption reveals distinct intraday patterns. During the early morning hours (00:00–05:00), consumption remains relatively low, reflecting minimal residential and industrial activity. A noticeable increase is observed starting around 06:00, corresponding to the beginning of daily operations in households and workplaces. Throughout the midday hours (10:00–17:00), consumption levels remain relatively stable, coinciding with typical working periods. In the evening (18:00–21:00), there is a secondary rise in energy usage, likely driven by residential demand as people return home, followed by a gradual decline during the late evening and nighttime.

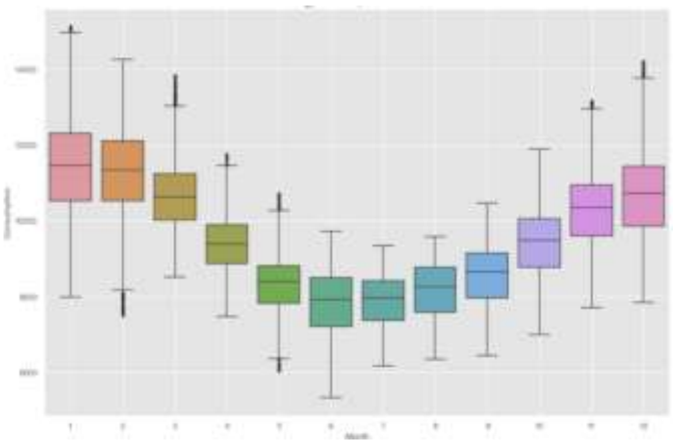

Fig -2: Energy consumption VS Month

Monthly patterns of energy consumption demonstrate a pronounced seasonality. Consumption peaks during the winter months, particularly between November and February, which can be attributed to increased heating requirements during colder weather. In contrast, energy usage declines significantly during the summer months of June through August, when milder temperatures and extended daylight hours reduce heating and lighting demands. Transitional months such as March, April, September, and October display moderate consumption values, reflecting the gradual shift between winter and summer energy needs.

A broader examination of yearly energy consumption from 2016 to 2021 reveals a consistent seasonal pattern, yet certain deviations are observed. While most years exhibit stable annual energy demand cycles, a noticeable decline is observed in 2020, aligning with the global impact of the COVID-19 pandemic. The reduction in energy consumption during 2020 is likely attributable to decreased industrial activity, remote working policies, and reduced public mobility, leading to lower overall energy demand. In 2021, a partial recovery in energy consumption can be seen, suggesting a rebound in economic and social activities.

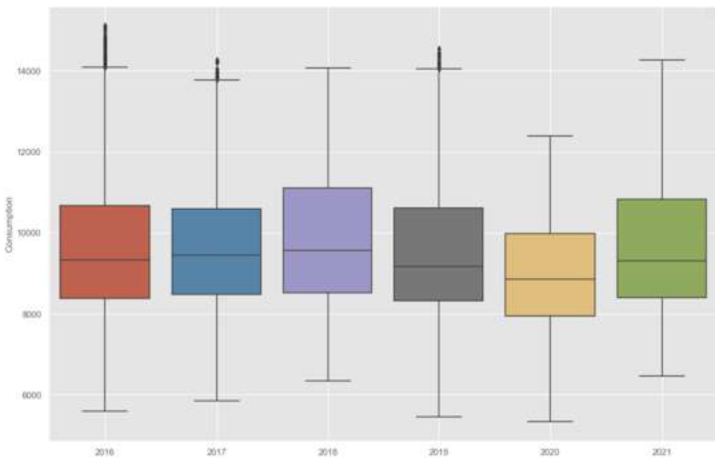

Fig -3: Energy consumption VS Year

In conclusion, the analysis of energy consumption trends across different time frames highlights clear seasonal patterns, with consumption peaking during winter months and dipping in summer. The hourly data reveals a distinct daily cycle, with higher consumption during peak hours in the evening, while the monthly breakdown demonstrates the prominence of winter as the primary driver of increased demand. The yearly analysis further underscores these seasonal fluctuations, with a noticeable dip in 2020, likely due to the economic effects of the COVID-19 pandemic. These insights provide a strong foundation for modeling energy consumption patterns, offering valuable context for forecasting future energy demand.

5. Predictions

The LSTM model for predicting daily energy consumption was built using down-sampled data from hourly to daily frequency, reducing the dataset from 52,774 rows to 2,184. To prepare the data for the LSTM model, normalization was performed using the Min-Max Scaler function, ensuring the values were scaled between 0 and 1. The dataset was split into training (80%), testing (20%), and validation sets (20% of the training data). The LSTM model was reshaped into a 3D array with time steps of 100 and features representing consumption.

A stacked LSTM model with four hidden layers, each containing 50 units, was created. A dropout layer was added to prevent overfitting, while the Adam optimizer and RMSE were used to evaluate the model's performance. The model was trained for 60 epochs with a batch size of 20, and during training, the loss values were monitored for both training and validation data to prevent overfitting.

The trained model's predictions were then tested on the training, validation, and test datasets, showing the model's ability to predict consumption values effectively. Figures depicting predictions on these datasets further illustrate the model's accuracy and its application to future consumption forecasts.

This section highlights the predictive power of the LSTM model, providing valuable insights for future energy consumption trends based on historical data.

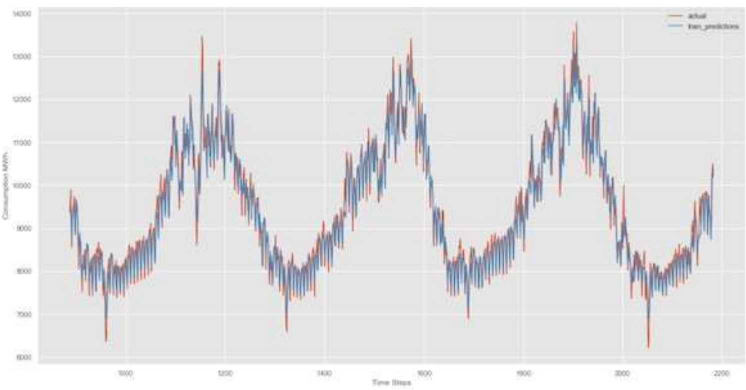

Fig -4: Actual vs train predictions

6. CONCLUSIONS

This study focused on analyzing and forecasting Finland's energy consumption using Long Short-Term Memory (LSTM) networks. Exploratory data analysis revealed distinct temporal patterns: energy usage exhibited clear hourly peaks during early mornings and evenings, monthly seasonality with higher consumption during winter months and lower during summers, and a relatively stable yearly trend except for a decline in 2020, attributed to the COVID-19 pandemic. These observations underline the importance of temporal dependencies in modeling energy demand accurately.

To capture these patterns, an LSTM-based deep learning model was developed. The dataset was resampled from hourly to daily frequency, normalized, and structured into training, validation, and testing subsets. A stacked LSTM architecture, enhanced with dropout regularization and trained using the Adam optimizer, was employed. The model demonstrated strong predictive performance with low training and validation errors and no significant overfitting, validating its capability to learn complex temporal relationships within the data.

Overall, the integration of time series analysis with deep learning proved to be effective in forecasting energy consumption trends. The findings emphasize that leveraging seasonal, daily, and yearly patterns through advanced neural network models can significantly enhance the accuracy of energy demand forecasting, supporting better energy planning and sustainable resource management.

ACKNOWLEDGEMENT

We sincerely thank our guide Dr. Anupama C.G. for their invaluable support and guidance throughout this research. We are also grateful to SRM institute of science and technology for providing the necessary resources. Lastly, we thank our family and friends for their constant encouragement.

REFERENCES

1. Electricity sector in Finland. (2022, September 23). Retrieved from Wikipedia: https://en.wikipedia.org/wiki/Electricity-sector-in-Finland
2. Time series forecasting methods. (n.d.). Retrieved from InfluxData: https://www.influxdata.com/time-series-forecasting-methods/
3. Fingrid Datahub. (n.d.). Electricity consumption datasets. Retrieved from Fingrid: https://data.fingrid.fi/en/datasets
4. What is time series data? (2022). Retrieved from https://www.influxdata.com
5. Brownlee, J. (2016). Deep Learning for Time Series Forecasting: Predict the Future with MLPs, CNNs and

LSTMs in Python. Machine Learning Mastery. https://machinelearningmastery.com

6. Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep Learning. MIT Press. Retrieved from https://www.deeplearningbook.org/
7. Statista. (2023). Electricity consumption in Finland. Retrieved from https://www.statista.com/statistics/602912/finland-electricity-consumption/