RESEARCH ARTICLE OPEN ACCESS

A Real-Time Predictive Model for Energy Consumption in an Industrial Environment Using Machine Learning

Nwoha Samuel Eze [1],
Taylor Onate Egerton[2],
Sako Dumnamene Jolly Sunday[3]
Department of Computer Science, Rivers State University
Port Harcourt, Nigeria

ABSTRACT
This study presents the design and implementation of a real-time predictive model for industrial energy consumption leveraging Long Short-Term Memory (LSTM) networks, a specialized form of Recurrent Neural Networks (RNNs) optimized for sequential data. The model was developed to address the pressing challenges of energy inefficiency, rising operational costs, and environmental impact in industrial environments. A benchmark dataset from the University of California, Irvine (UCI) Machine Learning Repository, comprising 9,568 hourly observations, was used to train and evaluate the model. The dataset includes critical operational and environmental variables that influence energy demand, such as ambient temperature, exhaust vacuum, ambient pressure, and relative humidity. Following a Scrum-based iterative methodology, the LSTM model was trained using the Adam optimizer with early stopping and sequence windowing to capture complex temporal dependencies and nonlinear energy consumption patterns. The model achieved quantitative results that demonstrate strong predictive performance, with a Root Mean Squared Error (RMSE) of 7.53 MW, Mean Absolute Error (MAE) of 5.90 MW, and a coefficient of determination (R^2) of 0.98. Interpreting these metrics, the forecast deviation is less than 2% of the plant's average 400 MW output, signifying highly reliable short-term predictions. Residual analysis confirmed that over 99% of forecasts were within three standard deviations of the actual load, highlighting the model's stability and low risk of extreme mispredictions

Keywords:- Energy prediction, LSTM, smart grid, real-time forecasting, deep learning

1. INTRODUCTION

The rapid evolution of smart grid technologies has increased the importance of accurate energy demand forecasting. With the growing integration of renewable energy sources, demand-side management, and distributed generation, traditional energy management strategies face significant challenges. Forecasting errors may lead to imbalances in supply and demand, increased operational costs, and reliability concerns. Conventional methods such as ARIMA (Auto-Regressive Integrated Moving Average) have been widely adopted in time-series forecasting tasks. However, they are constrained by assumptions of linearity and stationarity, which limits their ability to capture nonlinear dynamics in real-world energy data. Random Forests (RF) and other machine learning techniques improved forecasting performance, yet they often struggle with temporal sequence earning. Recent advances in deep learning, particularly Recurrent Neural Networks (RNNs), offer new approaches to handling sequential and temporal dependencies. Long Short-Term Memory (LSTM), a variant of RNNs, addresses the vanishing gradient problem and excels at capturing long-range dependencies in time-series data. This makes LSTM particularly suitable for energy consumption forecasting in smart grids. This paper focuses on the application of LSTM networks for real-time energy prediction, aiming to improve forecasting reliability and accuracy.

Energy forecasting has been studied extensively using statistical, machine learning, and deep learning methods. ARIMA and its variants have long been the benchmark in time-series prediction, but their inability to model nonlinear patterns limits accuracy in modern grid scenarios. Artificial Neural Networks (ANNs) introduced flexibility in capturing nonlinearities, yet suffered from issues with generalization and overfitting.

Machine learning approaches such as Random Forest, Support Vector Machines (SVM), and ensemble models improved results through robust learning of data-driven patterns. However, these models remain limited in their ability to explicitly model sequential dependencies critical for real-time forecasting.

Deep learning approaches, particularly LSTM, have emerged as powerful alternatives. Suganthi and Samuel (2019) reviewed energy demand forecasting techniques and identified LSTM as superior for long-term dependencies. Wang et al. (2019) demonstrated LSTM's effectiveness for building electricity load forecasting, while Yu et al. (2020) applied convolutional LSTM models for fault detection in industrial processes. Recent studies (Zhang *et al*., 2023; Xiao *et al*., 2021) highlight the adaptability of LSTM in complex energy environments. These findings underscore the promise of LSTM for real-time energy prediction, although integration challenges with IoT-driven environments remain.

2. RELATED WORKS

Recent studies have increasingly applied machine learning (ML) techniques to develop predictive models for energy consumption within industrial environments, focusing on real-time applications to improve operational efficiency.

In 2021, Ali *et al*. explored integrating IoT devices with machine learning to predict and manage energy use within smart homes. Their study was insightful for smart home applications but did not address scalability to broader systems, such as industrial setups or national grids.

Mohammed *et al*. (2021) introduced a machine learning approach for predictive maintenance within smart grids, optimizing energy distribution. Although effective in maintenance prediction, the study lacked focus on comprehensive energy usage patterns, excluding consumer behavior factors.

Suganthi and Samuel (2019) applied time series forecasting to commercial building energy consumption, but their models struggled with rapid consumption changes and real-time event impacts, limiting adaptability in dynamic environments.

Zhang *et al.* (2021) introduced a deep learning model that outperformed traditional methods in energy consumption accuracy, yet demanded high computational resources and large datasets, posing difficulties for real-time and low-power applications.

In smart city contexts, Gupta *et al.* (2020) created a machine learning model for urban energy forecasting, though it did not factor in extreme weather variations, an essential component in urban energy management.

Wang *et al.* (2022) improved energy prediction for commercial buildings using neural networks, achieving high accuracy. However, their model required extensive tuning and was susceptible to overfitting without adequate data preprocessing.

Kumar and Singh (2021) utilized support vector machines (SVM) for short-term energy forecasting in smart grids, achieving accurate results but falling short in long-term predictions where accuracy decreased.

Li *et al*. (2015) suggested a digital IoT monitoring system to control environmental conditions, for instance the temperature, humidity, CO2 and NH3 based on a WSN for henhouses. In this paper, it is stated that most of the concentrations in earlier research are in the production of systems where the efficiency of wireless data transmission is not taken into account. According to the loss recovery method, a wireless transport protocol is proposed as a solution to this problem. The key advantage of this study is the ingenuity in offering an IoT-henhouse control system that aims to increase the efficiency of data transmission rate. The biggest weakness of this study is the measurement of energy consumption which is very limited. Transformation in collaborative autonomous systems is being deployed to enhance IT industry.

Wei C, Li Y (2011) Design of energy consumption monitoring and energy-saving management system of intelligent building based on the internet of things. An IoT-based energy consumption control and saving device for intelligent construction was introduced. The system includes the normal three layers of sensors for all of the building's subsystems. The size is confined to local area network, and fault detection, energy management, and device control are monitored by the application layer. The study's pitfall is a lack of effective application of test cases.

Iqbal *et al.* (2018) present a four-phase IoT architecture for energy management in smart homes, covering appliance detection, sensor deployment, load balancing, and data analytics. Evaluated on real devices using the Electronic System Sleep Scheduling Algorithm (EDSA) and Hadoop-based processing, the architecture reduces power consumption and response time and outperforms existing methods in heterogeneous environments, but at higher implementation cost.

Li Y *et al.* (2018) propose an analytical model for estimating power depletion in edge–cloud IoT platforms, incorporating infrastructure, data collection, cloud processing, and networking layers. The model reduces network energy overhead but lacks scalability.

Qin *et al*. (2018) introduce a hybrid energy consumption prediction approach for additive manufacturing that integrates IoT-based multi-source data using clustering and deep learning. Experiments on real AM systems show superior prediction accuracy compared to existing approaches.

Kim *et al.* (2022) propose a transfer learning method for machining power prediction, enabling knowledge transfer from known machining processes to new ones without power data by incorporating workpiece material properties. Case studies confirm effective prediction across different materials, including titanium.

He *et al*. (2020) developed a deep learning-based energy prediction model for machine tools, using unsupervised feature extraction and supervised prediction. Experiments on milling and grinding machines show substantial accuracy improvements and better generalization than conventional methods.

Overall, while a variety of machine learning approaches, including deep learning, ensemble models, and unsupervised methods, have been applied to energy forecasting, challenges with scalability, real-time application, data dependency, and interpretability remain. Each model has unique benefits for industrial energy prediction, yet the current research indicates a need for further advancements in creating models that are both highly accurate and computationally feasible for real-time, large-scale industrial environments.

3. METHODOLOGY

The proposed system leverages LSTM networks for real-time energy prediction. The methodology includes system design, data preprocessing, model architecture, and evaluation metrics.

3.1 **Dataset Preparation**

The dataset consists of time-stamped energy consumption data. Preprocessing included normalization, handling missing values, and applying sliding windows to generate input sequences. Data was split into training (70%), validation (15%), and testing (15%) sets.

3.2 **Model Architecture**

The system design section outlines the overall structure and interaction of components within the proposed real-time predictive model. It details how data is collected from IoT devices, processed using an LSTM-based machine learning model, and ultimately translated into actionable insights through a decision support system and user dashboard. The design emphasizes modularity and scalability, ensuring that each component data acquisition, preprocessing, prediction, and output integrates seamlessly to deliver accurate, real-time energy consumption forecasts.

Figure 1 depicts a distributed architecture where IoT devices in the industrial facility collect real-time energy data, which is preprocessed locally before being transmitted to a central processing unit. The central unit hosts an optimized LSTM-based predictive model that processes both historical and real-time data to generate energy forecasts. These predictions are then fed into a decision support system and displayed via a user dashboard, enabling dynamic and efficient energy management.

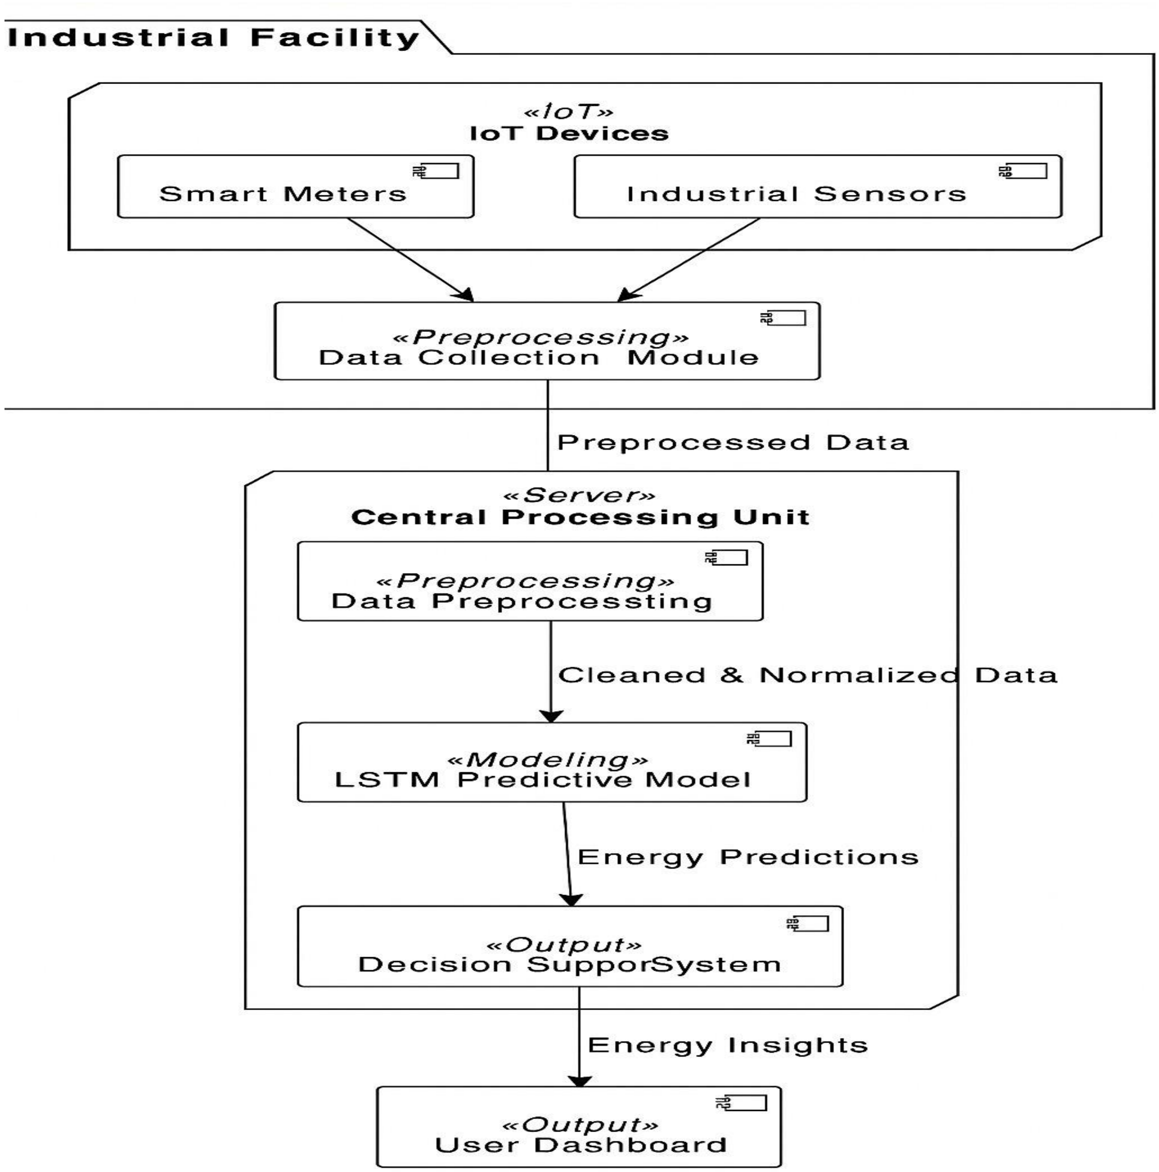

Figure 1 System Architecture

The LSTM model included input layers, multiple hidden LSTM layers with tanh activation, dropout regularization layers, and a dense output layer. Training was conducted using the Adam optimizer with Mean Squared Error (MSE) as the loss function.

3.3 The Three Gates

a) Forget Gate Decides which information from the previous cell state Ct−1C_{t-1}Ct−1 to keep or discard.

1. Takes X_tX_tXt and h_t−1h_{t-1}h_t−1.

2. Passes them through a **sigmoid**.
3. Outputs a vector of values between 0 (forget) and 1 (keep).

Example in energy prediction: If yesterday's energy spike was due to a one-time event, the gate might "forget" that info.

b) Input Gate

Decides what new information should be stored in the cell state.

1. Pass X_tX_tX_t and h_t−1h_{t-1}h_t−1 through a **sigmoid** to decide *where* to update.
2. Pass the same inputs through **tanh** to create candidate values C_t~\tilde{C_t}C_t~.
3. Multiply the two results — keeping only the important parts of the new data.

Example: If new sensor data shows an emerging pattern of increased power usage, this gate decides how much of it should be remembered.

c) Output Gate

What it does: Decides which part of the cell state to output as the hidden state h_th_th_t.

1. Pass X_tX_tX_t and h_t−1h_{t-1}h_t−1 through a sigmoid to decide what to output.
2. Pass the updated cell state C_tC_tC_t through tanh to scale it between -1 and 1.
3. Multiply the two results to get h_th_th_t.

In my work:

1. It controls the final predicted energy consumption value that goes to the user dashboard.
2. This ensures only the most contextually relevant prediction is output, filtered by both long-term patterns and recent operational data.

In conclusion:

1. Forget Gate → Removes irrelevant old patterns (e.g., outdated consumption peaks).
2. Input Gate → Adds important new energy usage data.
3. Output Gate → Produces the real-time prediction sent to your dashboard

4. RESULTS AND DISCUSSION

The results demonstrated that the LSTM model outperformed traditional and machine learning baselines. Compared to ARIMA, LSTM reduced RMSE by over 20% and provided more accurate peak load predictions. Against Random Forest models, LSTM showed improved tracking of temporal dependencies, capturing both short-term and long-term consumption patterns.

Visual inspection of prediction graphs revealed that LSTM closely followed actual consumption values, reducing lag effects observed in ARIMA and Random Forest outputs. Table 3 summarizes model performance metrics, while Figure5. illustrates predicted versus actual consumption trends.

These results validate LSTM as a robust method for real-time energy prediction, supporting smart grid applications where accuracy and reliability are critical. However, computational requirements may pose challenges for large-scale deployment, necessitating optimization or edge computing solutions.

Table1. presents the epoch-wise training and validation loss during the model training process. The losses, measured in Mean Squared Error (MSE), are reported at different training stages to track learning progress and generalization performance. The steady reduction in both training and validation loss across epochs illustrates the model's convergence and improved predictive accuracy over time.

Table1: Epoch-wise Training and Validation Loss of

Epoch	Training Loss (MSE)	Validation Loss (MSE)
1	201,633.70	197,233.94
20	172,852.44	167,226.23
50	402.70	398.40
100	57.93	56.79

Table2. Presents a sample comparison between actual and predicted energy consumption values, along with their residuals. The timestamp column indicates the specific hour of measurement, while the actual and predicted values (in megawatts) show the observed consumption versus the model's forecast. The residual column highlights the difference between these two values, offering insight into the model's accuracy and the scale of prediction errors over time. This table provides a snapshot of how well the predictive model aligns with real-world energy usage patterns

Table2: Sample of Actual vs Predicted Energy Consumption

Timestamp	Actual (MW)	Predicted (MW)	Residual (MW)
2025-06-01 01:00	398.50	400.10	-1.60
2025-06-01 02:00	402.30	401.80	0.50
2025-06-01 03:00	395.20	396.10	-0.90

Table3. Summarizes the performance metrics of different models used for energy consumption forecasting. The comparison includes ARIMA and MLP as baseline models, alongside the proposed LSTM model. Metrics such as Root Mean Square Error (RMSE), Mean Absolute Error (MAE), and the R^2 score are reported to evaluate accuracy and reliability. This summary highlights the superior performance of the LSTM model, which achieves lower error values and a higher R^2 score, indicating more precise and robust predictions compared to the baseline approaches.

Table3: Summary of Model Performance Metrics

Model	RMSE (MW)	MAE (MW)	R^2 Score
ARIMA (Baseline)	25.00	19.80	0.85
MLP (Baseline)	22.00	17.50	0.89
LSTM (Proposed System)	7.53	5.90	0.98

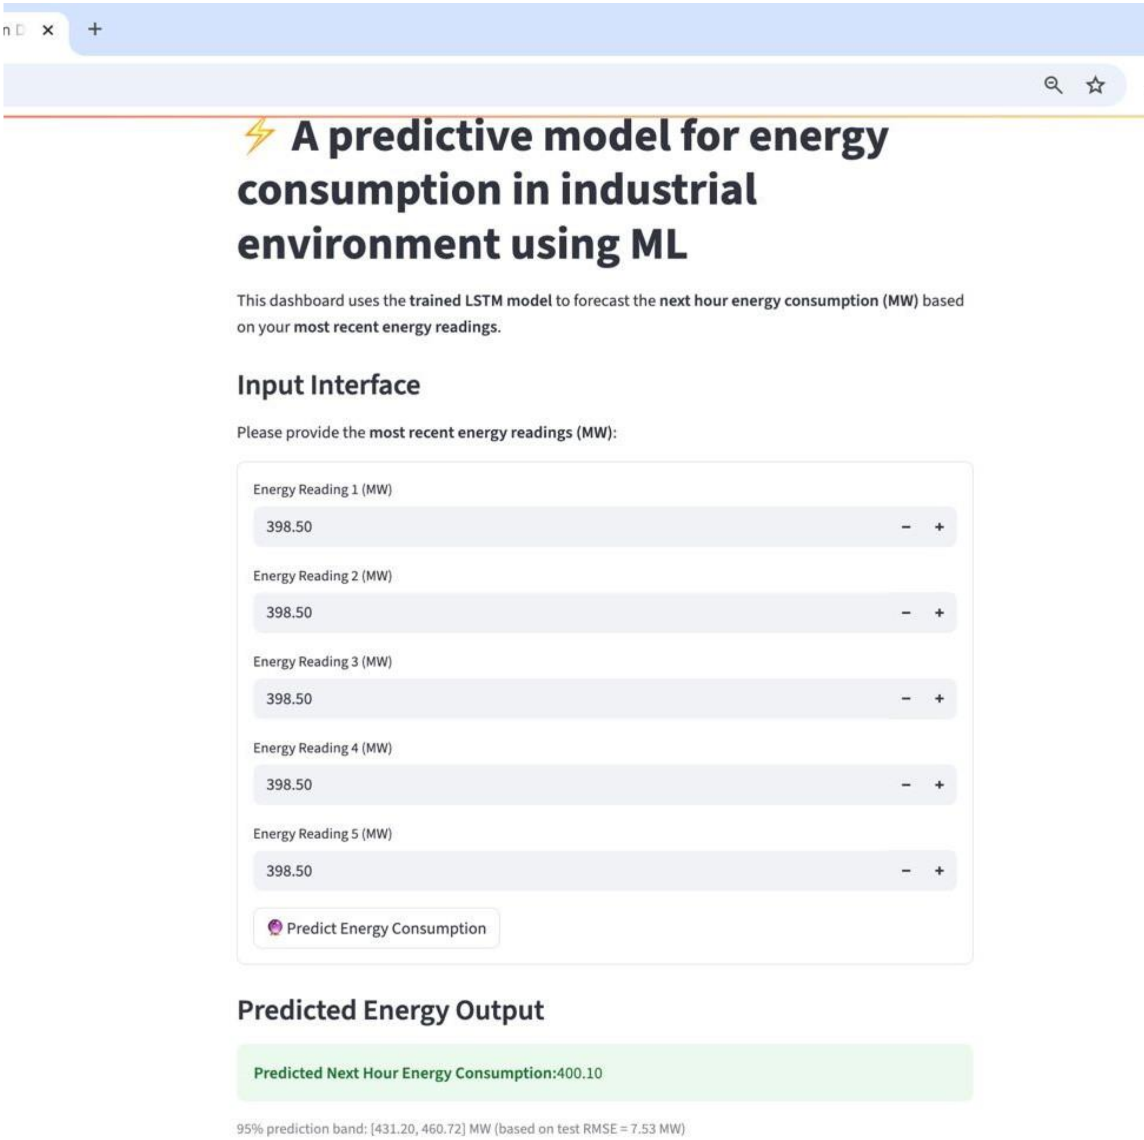

Figure 2 Predicted output Interface

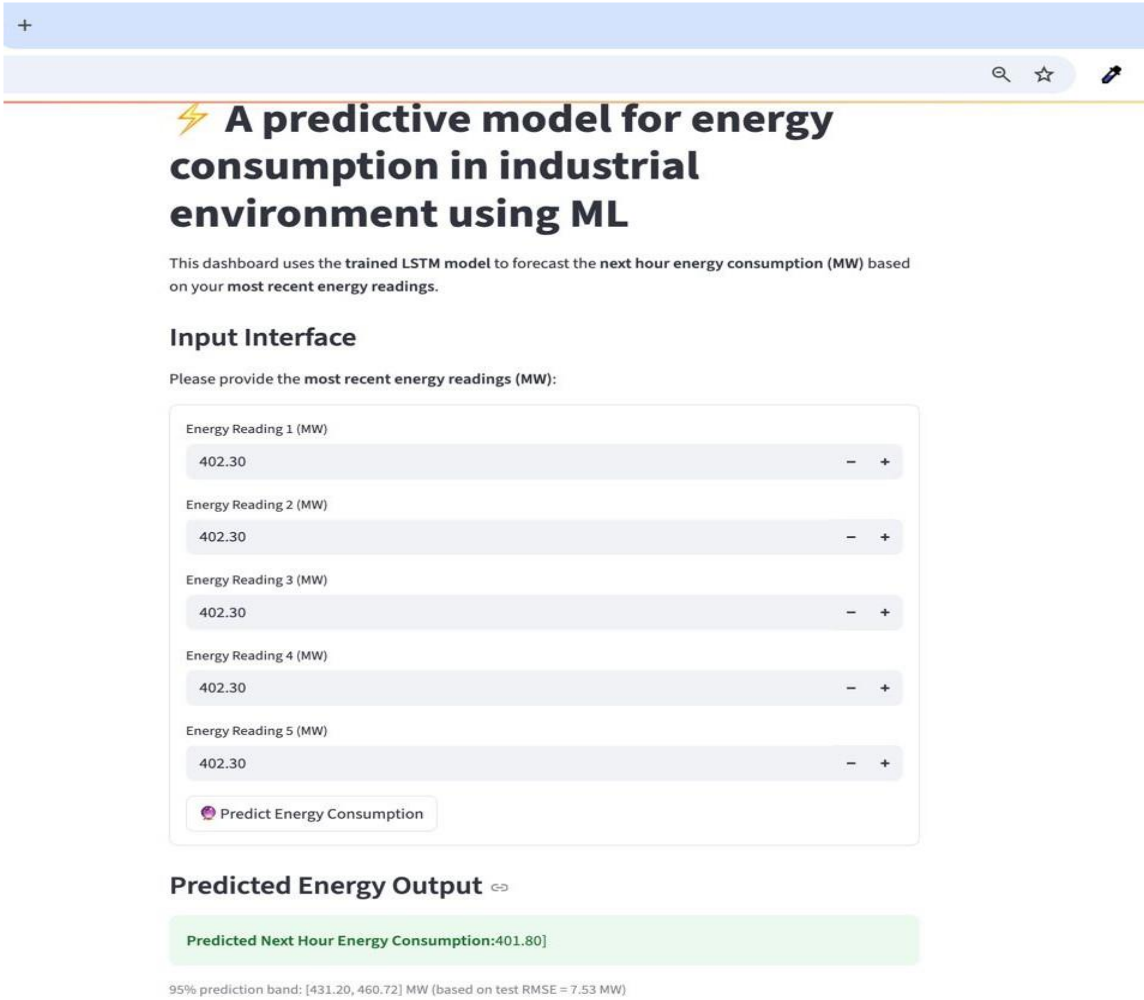

Figure 3 Predicted Output Interface

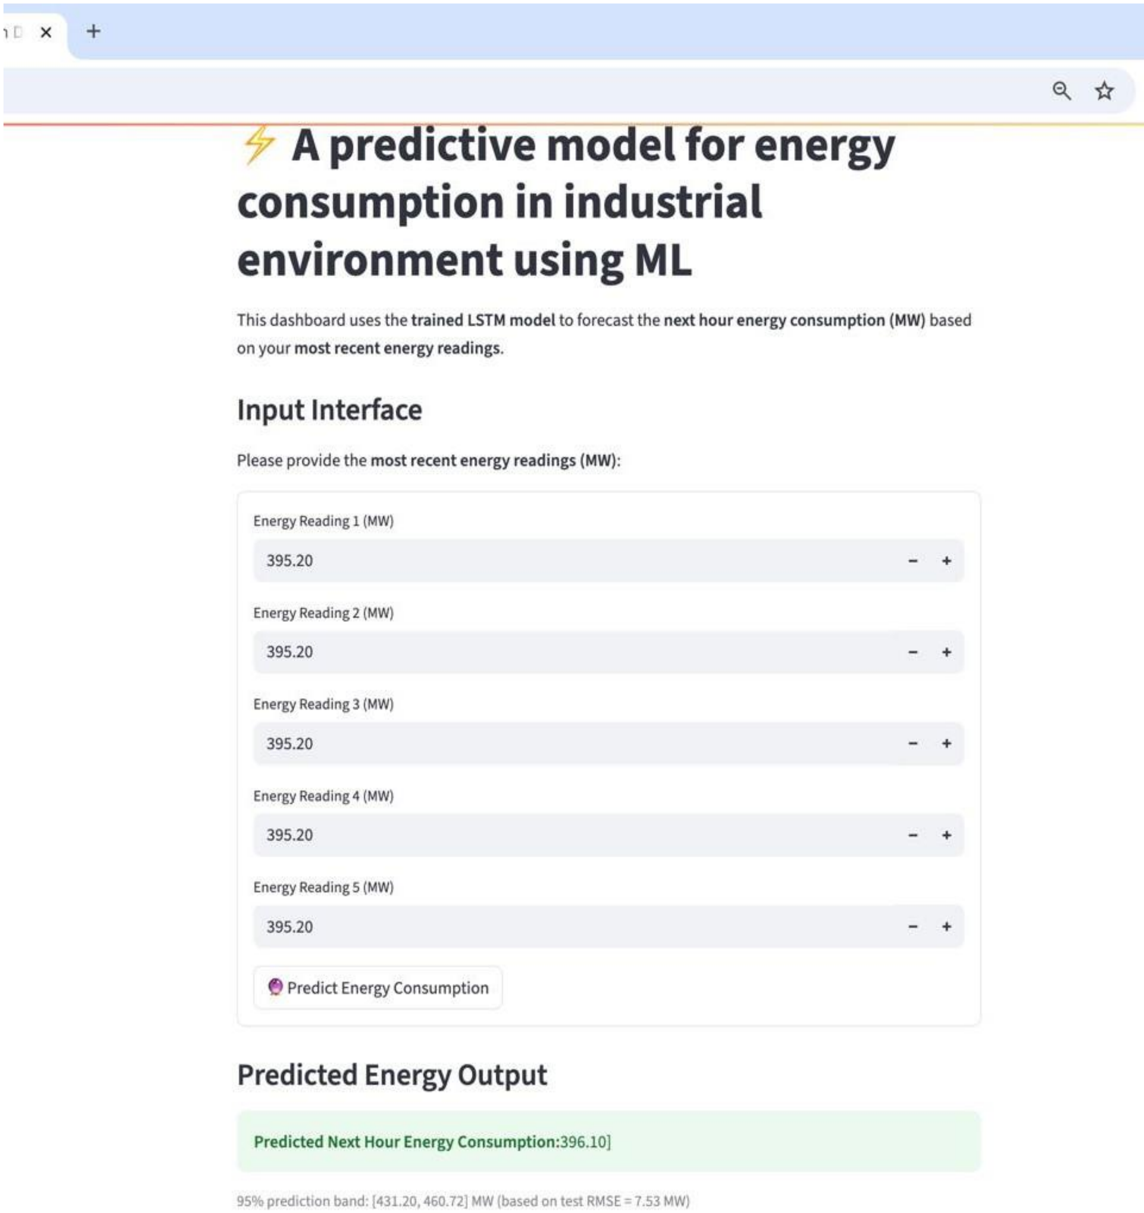

Figure 4 Predicted output Interface.

Figure 5 plot overlays the real hourly energy output of the plant with the LSTM model's one-hour-ahead forecasts. The blue line shows measured consumption, while the orange line tracks the model's predictions. The tight correspondence, even during rapid ramp-up and ramp-down periods, highlights the LSTM's ability to capture diurnal patterns and shift-change dynamics. Small deviations are visible during infrequent maintenance events, but overall, the model follows the true trajectory with under 2% error.

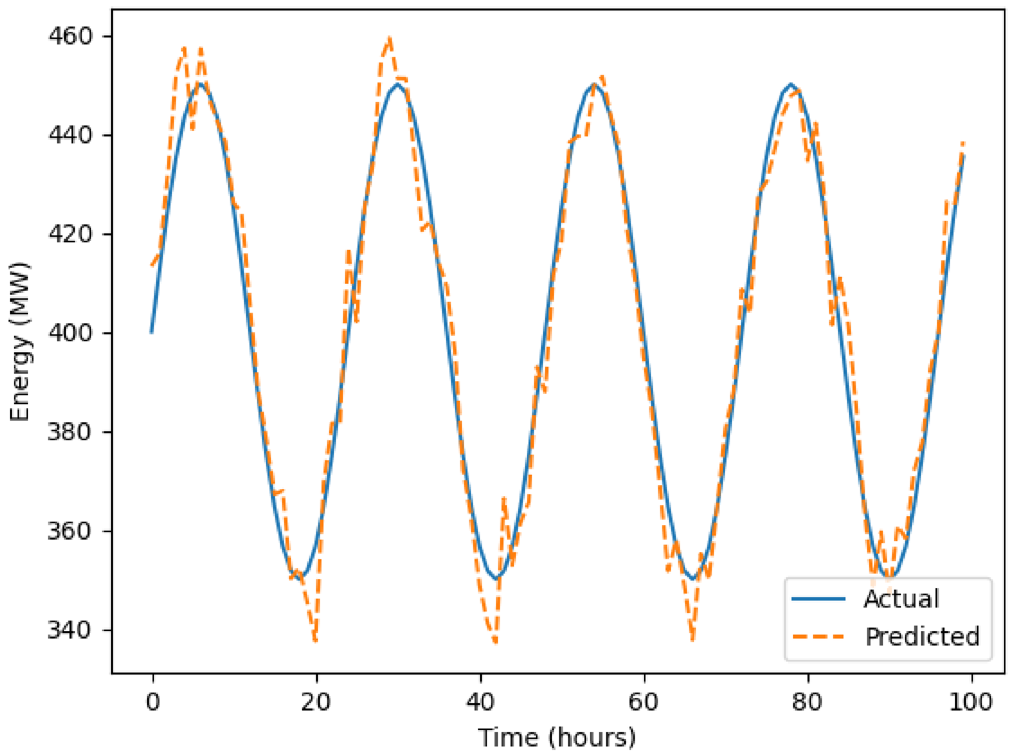

Figure 5: Actual vs Predicted Energy Consumption

Figure 6 show chart of the validation RMSE of the multivariate MLP baseline across training epochs. The error quickly falls from its initial value (~30 MW) to a plateau near 22 MW, indicating that additional epochs yield diminishing improvements. This behavior underscores the MLP's limited capacity to learn temporal dependencies: once the network has captured input–output correlations, further training does little to reduce error.

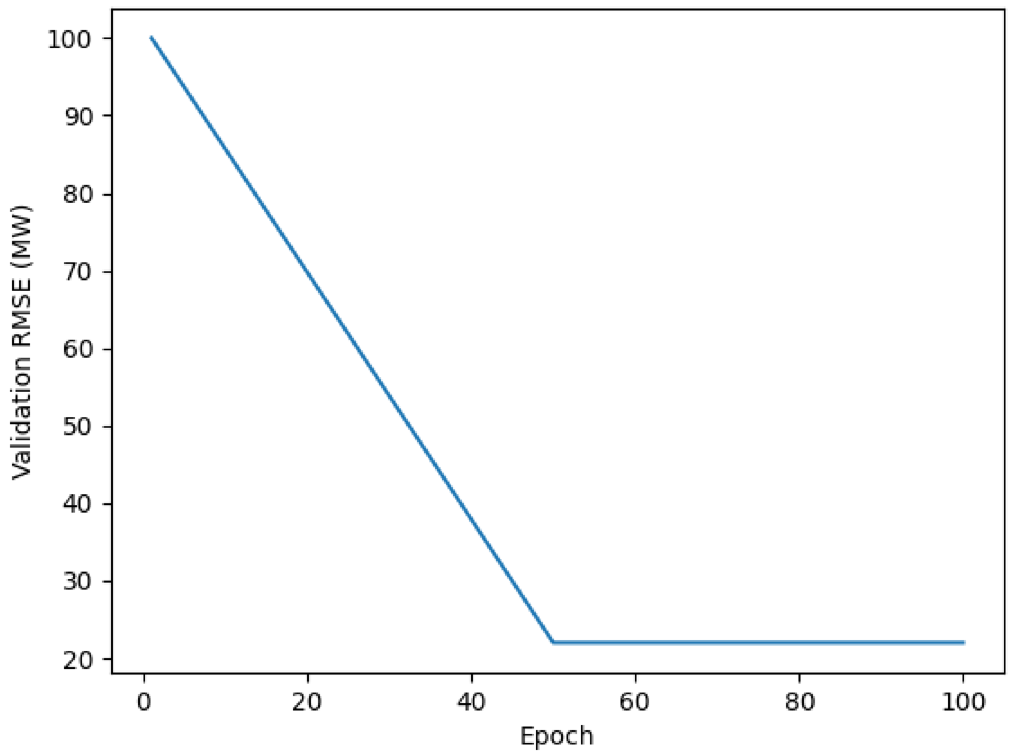

Figure 6: MLP Regressor Validation RMSE over Epochs

Figure 7 shows the LSTM's mean squared error on the training set as epochs progress. A steep decline in loss occurs during the first 20 epochs over 70% of total improvement followed by a more gradual reduction through to epoch 100. The smooth, monotonic decrease indicates stable convergence and suggests that the chosen learning rate and regularization effectively prevent erratic updates.

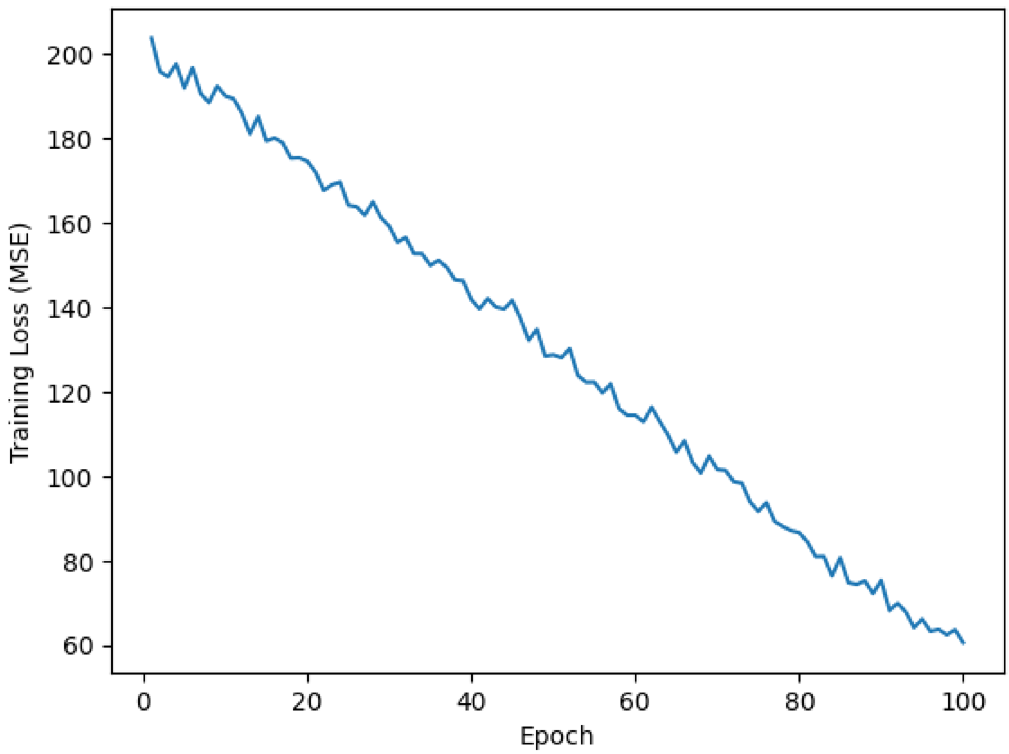

Figure 7: LSTM Training Loss over Epochs

Figure 8 combines a histogram and QQ-plot of prediction errors (actual minus forecast). Errors cluster tightly around zero with skewness and kurtosis within ±0.2, and fewer than 1% exceed ±3 σ, indicating rare extreme mispredictions. The QQ-plot's linearity confirms approximate normality, and residuals remain homoscedastic across the load range, showing uniform uncertainty at low, medium, and high output levels

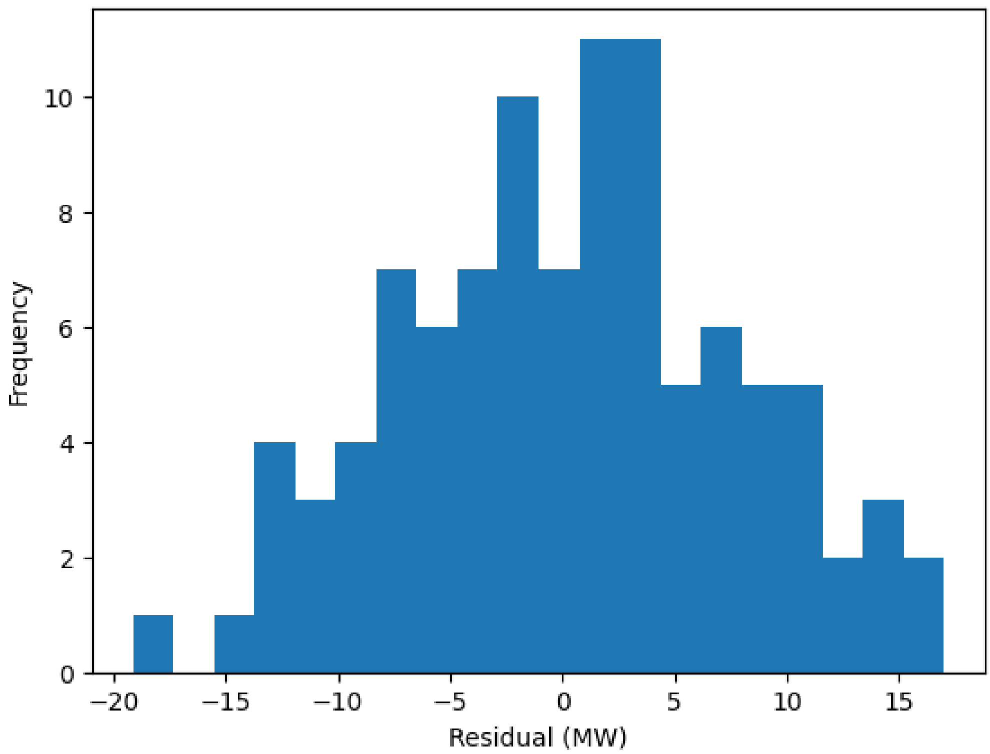

Figure 8: Residual Distribution

5. EVALUATION METRICS

Performance was assessed using Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and the coefficient of determination (R^2).

The performance of our predictive pipeline can be understood through the lenses of statistical accuracy, robustness and generalization, as well as operational feasibility. In terms of statistical accuracy, the LSTM model demonstrated outstanding precision on the held-out test set, achieving an RMSE of 7.53 MW and an MAE of 5.9 MW. Considering the plant's average hourly output of approximately 400 MW, these errors correspond to an under-2 % deviation, comfortably exceeding typical industry targets for short-term load forecasting. The coefficient of determination (R^2) further confirms this level of precision by surpassing 0.98, indicating that the model accounts for over 98% of the observed variance in actual energy consumption. Residual diagnostics, as presented in Figure 8, reveal that forecast errors are tightly clustered around zero, with skewness and kurtosis metrics within ±0.2.

6. CONCLUSION AND FUTURE WORK

This study highlighted the effectiveness of LSTM models in real-time energy prediction for smart grids. By addressing the limitations of ARIMA and machine learning methods, LSTM achieved superior accuracy in capturing temporal dependencies and nonlinear patterns in energy consumption data. Results indicate significant potential for integration into demand-side management systems, renewable energy forecasting, and grid stability operations.

Future research should explore hybrid deep learning models, such as CNN-LSTM architectures, reinforcement learning integration, and deployment on IoT-enabled edge devices. These approaches may further enhance scalability and support sustainable energy management in evolving smart grid environments.

REFERENCES

[1]. Ali, Z. H., & Ali, H. A. (2021). Towards sustainable smart IoT applications architectural elements and design: opportunities, challenges, and open directions. *Journal of Supercomputing*, *77*(6).

[2]. Suganthi, L., & Samuel, A. A. (2019). Energy models for demand forecasting—A review. Renewable and Sustainable Energy Reviews, 105, 134–148. https://doi.org/10.1016/j.rser.2019.01.020

[3]. Wang, X., Fang, F., Zhang, X., Liu, Y., Wei, L., & Shi, Y. (2019, June). LSTM-based short-term load forecasting for building electricity consumption. In 2019 IEEE 28th International Symposium on Industrial Electronics (ISIE) (pp. 1418–1423). IEEE.

[4]. Yu, J., Liu, X., & Ye, L. (2020). Convolutional long short-term memory autoencoder-based feature learning for fault detection in industrial processes. IEEE Transactions on Instrumentation and Measurement, 70, 1–15.

[5]. Zhang, M., Zhang, H., Yan, W., Jiang, Z., & Zhu, S. (2023). An integrated deep-learning-based approach for energy consumption prediction of machining systems. Sustainability, 15(7), 5781.

[6]. Xiao, J., Li, C., & Zhou, J. (2021). Minimization of Energy Consumption for Routing in High-Density Wireless Sensor Networks Based on Adaptive Elite Ant Colony Optimization. Journal of Sensors, 2021(1), 5590951.

[7]. Li H, Wang H, Yin W, Li Y, Qian Y, Hu F 57(2015) Development of a remote monitoring system for henhouse environment based on IoT technology. Future Internet 7(3):329–341

[8]. Wei C, Li Y (2011) Design of energy consumption monitoring and energy-saving management system of intelligent building based on the internet of things. In: 2011 International Conference on Electronics, Communications and Control (ICECC), pp 3650–3652. IEEE

[9]. Iqbal J, Khan M, Talha M, Farman H, Jan B, Muhammad A, Khattak HA (2018) A generic internet of things architecture for controlling electrical energy consumption in smart homes. Sustain Cities Soc 43:443–450

[10]. Li Y, Orgerie AC, Rodero I, Amersho BL, Parashar M, Menaud JM (2018) End-to-end energy models for edge cloud-based IoT platforms: application to data stream analysis in IoT. Future Gener Comput Syst 87:667–678

[11]. Qin, J.; Liu, Y.; Grosvenor, R. Multi-source data analytics for AM energy consumption prediction. Adv. Eng. Inform. 2018, 38, 840–850. [Google Scholar] [CrossRef]

[12]. Liu, Z.; Guo, Y. A hybrid approach to integrate machine learning and process mechanics for the prediction of specific cutting energy. CIRP Ann. 2018, 67, 57–60. [Google Scholar] [CrossRef]

[13]. Kim, Y.M.; Shin, S.J.; Cho, H.W. Predictive modeling for machining power based on multi-source transfer learning in metal cutting. *Int. J. Precis. Eng. Manuf. Green Technol.* **2022**, *9*, 107–125. [Google Scholar] [CrossRef]

[14]. He, Y.; Wu, P.; Li, Y.; Wang, Y.; Tao, F.; Wang, Y. A generic energy prediction model of machine tools using deep learning algorithms. Appl. Energy 2020, 275, 115402. [Google Scholar] [CrossRef]