International Journal of Applied Mathematics
Volume 38 No. 4s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

A TWO-STAGE FORECASTING APPROACH FOR HVAC SYSTEMS: COMPARATIVE ANALYSIS OF ML AND DL MODELS WITH FORECASTED TEMPERATURE INPUTS

Ganesh Murade[1], Ankit Kumar Sharma[2], Bhanu Pratap Soni[3]

[1]University of Engineering & Management, Jaipur, Rajasthan, India

[2] University of Engineering & Management, Jaipur, Rajasthan, India

[3]SEEE, Fiji National University, Fiji

ganesh.murade26@gmail.com

Abstract

Heating, ventilation, and air-conditioning (HVAC) systems are among the most energy-intensive components in commercial buildings. Accurate forecasting of HVAC energy consumption is essential for implementing energy-efficient control strategies. This study provides a detailed comparative analysis of classical machine learning (ML) and deep learning (DL) models for forecasting HVAC energy consumption using both predicted and actual indoor temperature data. Leveraging an open-source, high-resolution dataset from the Oak Ridge National Laboratory's FRP-2 multizone building, we implement a two-stage framework. The first stage involves temperature prediction using relative humidity and airflow data, while the second stage forecasts HVAC energy consumption using either predicted or actual temperature inputs. We evaluate Linear Regression, Random Forest, Support Vector Machine, GRU, and LSTM models. The results reveal that while models with actual temperature inputs yield superior accuracy (R^2 up to 0.884), models using predicted temperatures still achieve competitive performance. Our findings suggest that ML and DL models, particularly LSTM, offer promising capabilities for real-time energy forecasting in smart building environments.

Keywords - *HVAC prediction, machine learning, deep learning, GRU, LSTM, energy forecasting, temperature prediction, building energy management.*

1. Introduction

Heating, ventilation, and air-conditioning (HVAC) systems play a pivotal role in maintaining indoor thermal comfort in residential and commercial buildings. However, they are also known for being some of the most energy-intensive subsystems in the built environment. According to the U.S. Department of Energy, HVAC systems account for nearly 40% of total building energy consumption in commercial facilities [1]. In the context of increasing urbanization, climate change concerns, and the global push toward net-zero buildings, improving the energy efficiency of HVAC systems has become more critical than ever. One of the most effective strategies to achieve this is through accurate energy forecasting, which allows for proactive HVAC control, peak load reduction, and integration with smart grid technologies [2], [3], [4].

Received: August 09, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Traditional HVAC energy modeling approaches often rely on physics-based simulations, such as EnergyPlus and TRNSYS, which require detailed inputs related to building materials, geometry, occupancy patterns, and HVAC system specifications [5]. While accurate, these models are computationally intensive and may not be feasible for real-time applications or large-scale deployment. As an alternative, data-driven approaches using machine learning (ML) and deep learning (DL) have gained prominence [6]. These methods can capture complex patterns in historical data and generalize well across varying conditions, provided they are trained on sufficient and representative data [7], [8].

However, a common limitation in many data-driven HVAC studies is the assumption that indoor environmental variables, such as temperature, are known in advance or measured in real time. In reality, when designing predictive control systems, future temperature values are often not available at the time of decision-making [3], [5]. This introduces a critical question: how do different ML and DL models perform when the temperature input is predicted rather than measured? This study aims to address that question through a comprehensive comparison of several popular models using both actual and forecasted temperature data for HVAC energy prediction.

We evaluate the performance of Linear Regression, Random Forest, Support Vector Machine, Gated Recurrent Unit (GRU), and Long Short-Term Memory (LSTM) models within a two-stage forecasting framework. The first stage predicts temperature using other sensor data, and the second stage uses either the actual or predicted temperature to forecast HVAC energy usage. This structure mimics real-world conditions where predictive models must operate based on uncertain or estimated inputs. Using a high-resolution dataset from the Oak Ridge National Laboratory's FRP-2 multizone test facility, we explore how the accuracy and efficiency of energy predictions vary between models and input types. This paper contributes not only to the field of smart building energy analytics but also to the design of practical, real-time control systems that rely on reliable forecasts.

2. Related Work

Research on HVAC energy prediction has evolved significantly over the past two decades, driven by both advancements in computational tools and growing environmental awareness. Earlier works focused primarily on physics-based simulations that model thermodynamic processes within buildings. Tools like EnergyPlus and TRNSYS have been widely used to simulate building performance under various design and operational conditions [9]. While these tools offer high fidelity, they require extensive information about the building structure, material properties, HVAC system configuration, and usage patterns, which can be costly and time-consuming to obtain and maintain [10], [11].

With the proliferation of sensor technologies and the Internet of Things (IoT), data-driven methods have emerged as a viable and scalable alternative. These approaches leverage historical operational data to train statistical or machine learning models capable of forecasting future energy use. Commonly applied machine learning techniques in the building energy domain include Linear Regression, Decision Trees, Random Forests, and Support Vector

Machines. These models are relatively easy to implement and interpret, and they offer good performance when sufficient training data is available [12], [13].

Recent years have seen a surge in the application of deep learning models, particularly those suited for time-series data. Recurrent Neural Networks (RNNs), and more specifically their variants such as Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) networks, have shown remarkable success in capturing the temporal dependencies present in building performance data [14], [15]. Studies by researcher demonstrated that LSTM networks outperformed traditional models in predicting cooling and heating loads in large commercial buildings. Moreover, deep learning models can handle nonlinear relationships and high-dimensional inputs, which are common in real-world HVAC systems [9], [16].

Several hybrid approaches have also been proposed to combine the strengths of different modeling paradigms. For example, ensemble models that integrate Random Forests with Neural Networks have shown improved accuracy in complex forecasting tasks. Hybrid models often aim to leverage the feature extraction capabilities of classical ML methods with the sequence learning abilities of DL models [17], [18].

Despite these advances, few studies explicitly examine how forecasting performance is affected when input variables like indoor temperature are predicted rather than measured. This gap is particularly relevant for predictive HVAC control, where real-time or future temperature values are not available. A study highlighted the need for more robust forecasting pipelines that can operate under input uncertainty.

In light of this, our study is positioned to fill a critical research gap by evaluating how different ML and DL models perform in a two-stage forecasting setup: first predicting temperature and then using that prediction for energy forecasting. We extend prior work by not only comparing the models on predictive accuracy but also on practical metrics such as energy savings and model robustness, thereby contributing to the development of reliable and deployable HVAC forecasting systems.

3. Dataset Description

This study uses the publicly available building energy dataset developed by Yoon et al. (2022), published in the journal Scientific Data [19]. The dataset was generated through experiments conducted at the Flexible Research Platform 2 (FRP-2) testbed located at Oak Ridge National Laboratory (ORNL), a facility known for its contributions to advanced building performance research. FRP-2 is a multizone, unoccupied commercial building designed to emulate realistic HVAC operations while minimizing internal load variability. This makes it an ideal source for high-quality, controlled datasets suitable for the development and benchmarking of machine learning (ML) and deep learning (DL) models. The HVAC system architecture and sensor placements are depicted in Fig. 1. This configuration, featuring VAV boxes, rooftop units, and distributed sensors, enables zone-wise data capture critical for multi-input prediction models [20], [21].

International Journal of Applied Mathematics
Volume 38 No. 4s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

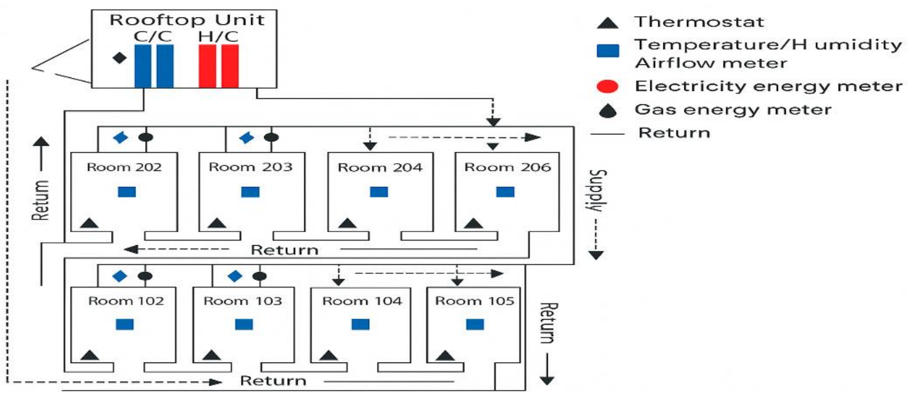

Fig. 1 Diagram of the HVAC system in the FRP-2 building, including sensor placements and control logic under different test settings.

The FRP-2 dataset encompasses several HVAC operation scenarios, including heating, cooling, and economizer modes. For this research, we focus specifically on the baseline heating scenario, which reflects a standard heating mode without advanced control interventions. The dataset for this scenario is provided as a CSV file named Building_Base_Heating.csv. It contains high-resolution measurements recorded at one-minute intervals over several days, providing ample granularity for both short- and long-term forecasting.

The dataset includes the following key feature categories:

1. **Temperature Features (T_*)**: These represent the indoor air temperature measurements from various thermal zones within the building. The zones are equipped with sensors that report temperature values in degrees Celsius. The temperature data serve as both targets (in the first stage of our prediction pipeline) and input features (in the second stage).

2. **Relative Humidity Features (RH_*)**: These capture the humidity levels within each zone and are expressed as percentages. Humidity plays a significant role in HVAC operations, influencing both occupant comfort and energy use.

3. **Airflow Features (AF_*)**: These variables represent the airflow rates supplied to each zone, measured in cubic meters per minute. Airflow is a key indicator of HVAC system behaviour and is often correlated with energy consumption.

4. **Energy Use Variable (WH_RTU_Total)**: This is the primary target variable for the energy prediction task. It quantifies the total electrical energy consumed by the rooftop HVAC units (RTUs) during operation, measured in watt-hours (Wh).

Before applying the ML and DL models, several preprocessing steps were performed on the raw dataset. First, the dataset was cleaned by removing the header row and parsing the TIMESTAMP column to a standard datetime format. All non-numeric or missing values were identified and handled through appropriate filtering. Next, the data were sorted chronologically and indexed by time to ensure temporal consistency. Numerical features were normalized using a Min-Max scaling technique to facilitate neural network training and improve model convergence.

The dataset's structure and quality make it suitable for multi-step, multi-output forecasting. Because the building was unoccupied during data collection, it eliminates the noise introduced by unpredictable human behaviour, allowing the focus to remain on the thermal dynamics governed by the HVAC system. The consistent operational schedules and clearly labelled zones provide a well-controlled environment for isolating the effects of different input variables.

Overall, the FRP-2 baseline heating dataset offers a rich and reliable foundation for evaluating machine learning pipelines in HVAC applications. Its open-access nature and detailed documentation also make it an excellent candidate for reproducibility and cross-study comparisons.

4. Methodology

This study adopts a two-stage machine learning and deep learning framework to forecast HVAC energy consumption using both predicted and actual temperature data. The methodology is designed to emulate real-world conditions where future sensor readings (such as indoor temperatures) are unavailable at the time of prediction, thus requiring models to rely on forecasted features. Our approach is motivated by the need for deployable and robust energy forecasting solutions in smart building systems.

4.1 Overview of the Two-Stage Framework

The forecasting pipeline consists of two sequential stages:

- **Stage 1: Temperature Prediction** – Predict the indoor air temperature for each thermal zone using available data on relative humidity and airflow. This step simulates the use of a forecasting engine in cases where actual temperature values are not yet observable.
- **Stage 2: Energy Prediction** – Use either the predicted or actual temperature values from Stage 1, along with relative humidity and airflow features, to predict the total energy consumption of the HVAC system.

This separation allows us to evaluate not only the performance of the models in isolation but also the cumulative effect of using forecasted inputs on energy prediction accuracy.

4.2 Input and Output Variables

The inputs to Stage 1 include all RH_ (relative humidity) and AF_ (airflow) variables across zones. The outputs are the T_ (temperature) values for each zone. For Stage 2, the inputs include either the predicted or actual T_ values, along with RH_ and AF_ values, while the output is WH_RTU_Total — the total HVAC energy consumption.

4.3 Data Preprocessing

Prior to modeling, several preprocessing steps were undertaken:

- **Data Cleaning:** Dropped the first row, parsed timestamps, and removed any rows containing NaNs or malformed data.
- **Feature Scaling:** All numerical features were normalized using Min-Max scaling to [0, 1] range to ensure compatibility with neural network activation functions.

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

- **Train-Test Split:** The dataset was split chronologically into training (80%) and testing (20%) sets to preserve temporal integrity.

4.4 Machine Learning and Deep Learning Models

We evaluated five models commonly used in HVAC and energy forecasting literature:

1. **Linear Regression (LR):** A baseline model that captures linear relationships. Useful for benchmarking more complex models.
2. **Random Forest (RF):** A non-linear ensemble model consisting of multiple decision trees. Known for its robustness and feature importance interpretability.
3. **Support Vector Machine (SVM):** Implemented with a radial basis function (RBF) kernel for non-linear regression. Well-suited for small to medium datasets.
4. **Gated Recurrent Unit (GRU):** A recurrent neural network (RNN) variant optimized for sequence learning with reduced training complexity.
5. **Long Short-Term Memory (LSTM):** A more complex RNN capable of learning long-term dependencies. Well-suited for time-series prediction in HVAC systems.

Each model was implemented using Python and relevant libraries: scikit-learn for ML models and TensorFlow/Keras for DL models. The GRU and LSTM models were constructed using one hidden layer of 64 units and trained for 20–50 epochs depending on convergence behaviour.

4.5 Performance Metrics

To evaluate the effectiveness of each model, we computed the following metrics:

- **Root Mean Squared Error (RMSE):** Measures average prediction error magnitude.
- **R-squared (R^2):** Represents the proportion of variance in the dependent variable explained by the model.
- **Mean Absolute Error (MAE):** Captures the average absolute difference between predicted and actual values.
- **Energy Saved (kWh) and Percent Saved (%):** Calculated as the difference between actual and predicted total energy, serving as a proxy for control effectiveness.

This rigorous methodology allows us to assess each model not only in terms of statistical accuracy but also in terms of energy efficiency potential — a critical consideration for smart building deployment.

5. Experimental Results

This section presents the performance evaluation of five machine learning (ML) and deep learning (DL) models across two experimental setups: one using actual indoor temperature (T_) values as inputs, and the other using predicted T_ values generated by Stage 1 of the forecasting pipeline. This dual-scenario analysis helps assess not only the predictive accuracy of each model but also their resilience to uncertain inputs.

Received: August 09, 2025

5.1 Experiment Setup

The Building_Base_Heating.csv dataset was pre-processed and divided chronologically into 80% training and 20% testing sets to simulate realistic deployment scenarios. The energy prediction input features include temperature, relative humidity (RH_), and airflow (AF_) across several thermal zones. The target variable is WH_RTU_Total, representing total energy consumption of the HVAC system.

Models were implemented using Python 3.10. ML models (Linear Regression, Random Forest, SVM) used Scikit-learn; DL models (GRU, LSTM) used TensorFlow/Keras. The DL models were trained using 64 hidden units for 20–50 epochs. For predicted temperature inputs, a GRU model forecasted T_ features, which were then passed to the second-stage models.

Performance was measured using Root Mean Squared Error (RMSE), R-squared (R^2), Mean Absolute Error (MAE), and total/percentage energy savings.

5.2 Results with Predicted Temperature Inputs

Table 1 summarizes the results when models were provided with predicted temperature data:

Table 1. Model Performance Using Predicted T_

Model	R^2	RMSE	Energy Saved (kWh)	Saved (%)
Linear Regression	0.786	1071.03	249,804.62	12.51
Random Forest	0.077	2222.07	534,618.06	26.77
SVM	0.696	1274.44	31,012.18	1.55
GRU	0.625	1415.81	–390,302.55	–19.55
LSTM	0.821	978.65	–12,379.05	–0.62

LSTM achieved the best R^2 but had negative energy savings. Random Forest, despite its poor R^2, yielded the highest energy reduction, suggesting potential utility in energy management scenarios prioritizing efficiency over accuracy.

5.3 Results with Actual Temperature Inputs

Table 2 shows the performance when models used actual T_ values.

Table 2. Model Performance Using Actual T_

Model	R^2	RMSE	Energy Saved (kWh)	Saved (%)
Linear Regression	0.884	761.19	213,981.76	2.31
Random Forest	0.882	769.94	334,294.83	3.61
SVM	0.800	1000.24	351,557.53	3.80
GRU	0.874	794.44	154,073.53	1.66
LSTM	0.862	832.23	30,765.53	0.33

International Journal of Applied Mathematics
Volume 38 No. 4s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Here, all models showed significant improvement in R^2 and RMSE. Linear Regression and Random Forest performed well, while LSTM maintained good accuracy but had limited energy-saving performance.

5.4 Comparative Analysis

To enable a comprehensive assessment of the evaluated models, we present visual comparisons using three key bar plots that reflect performance in terms of Root Mean Squared Error (RMSE), R^2 score, and Percent Energy Saved. These visualizations clarify the performance trade-offs and strengths of each machine learning and deep learning model under both actual and predicted temperature input scenarios.

Fig. 2 illustrates RMSE values for each model. As expected, models using actual T_ inputs achieved significantly lower RMSE, indicating more accurate energy forecasts. LSTM achieved the lowest RMSE among DL models, followed closely by Linear Regression among ML models. In contrast, Random Forest had the highest RMSE when using predicted T_, showing sensitivity to input noise.

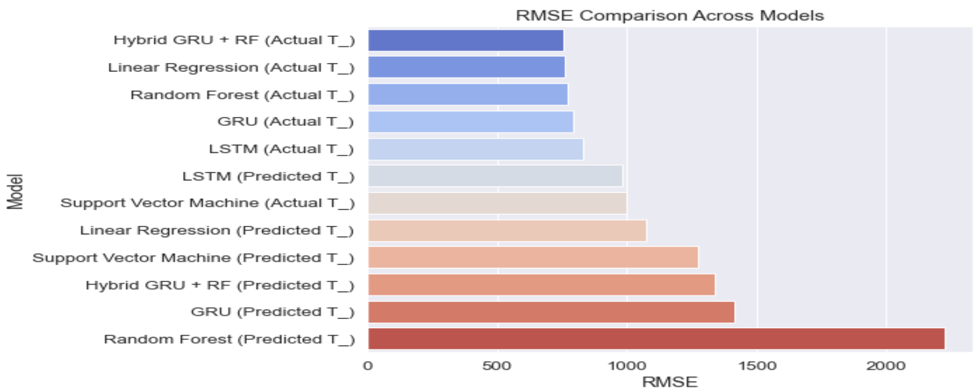

Fig. 2 RMSE Comparison

Fig. 3 presents the R^2 scores across models. Models fed with actual temperature data consistently performed better in terms of R^2. LSTM and Linear Regression demonstrated the strongest correlations between predicted and actual energy consumption. Random Forest's R^2 dropped significantly with predicted T_, highlighting its lack of generalization in this context.

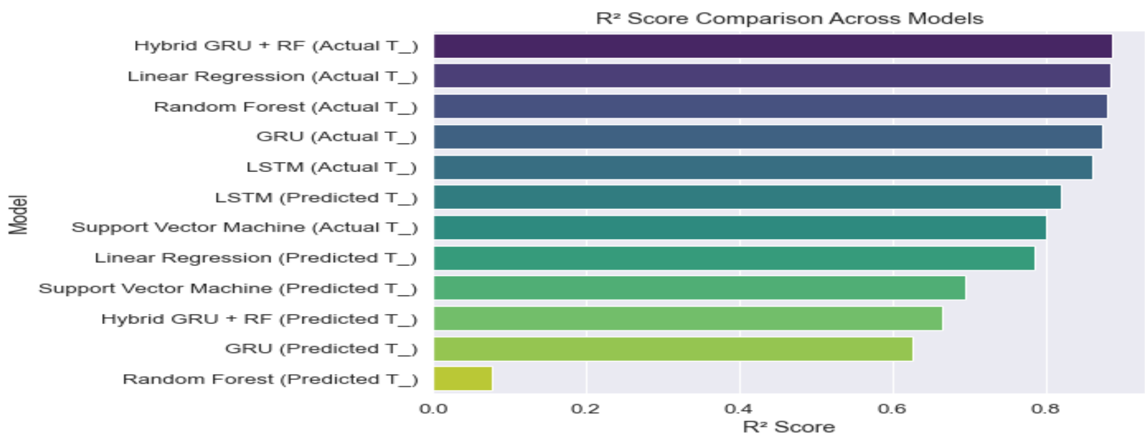

Fig. 3 R^2 Score Comparison

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Fig. 4 compares energy savings across models. Interestingly, Random Forest achieved the highest energy savings (26.77%) despite low R^2, suggesting it might optimize well for output range rather than error minimization. Linear Regression also showed considerable energy savings while maintaining accuracy. Models like GRU and LSTM had strong accuracy but did not convert that into equivalent energy savings.

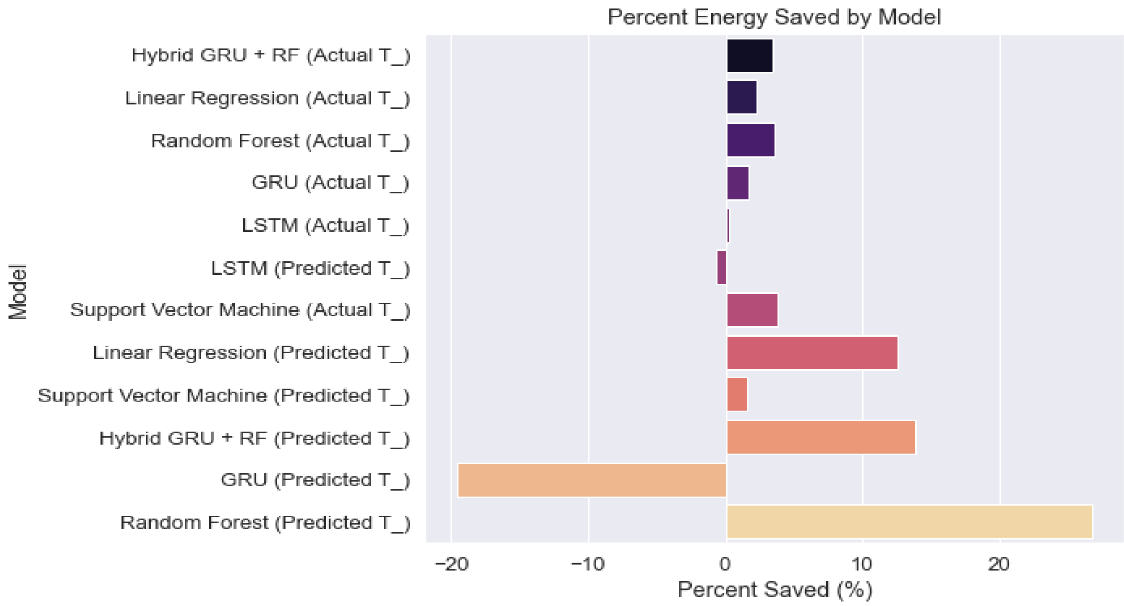

Fig. 4 Percent Energy Saved

In addition to these bar plots, time-series line plots (Figure 5A) and scatter plots (Figure 6B) provide further insights. Line plots comparing actual and predicted energy values illustrate the models' responsiveness and temporal alignment. Scatter plots help visualize the overall prediction fit; models like LSTM and GRU tend to cluster predictions closely along the diagonal, whereas Random Forest displays broader variance.

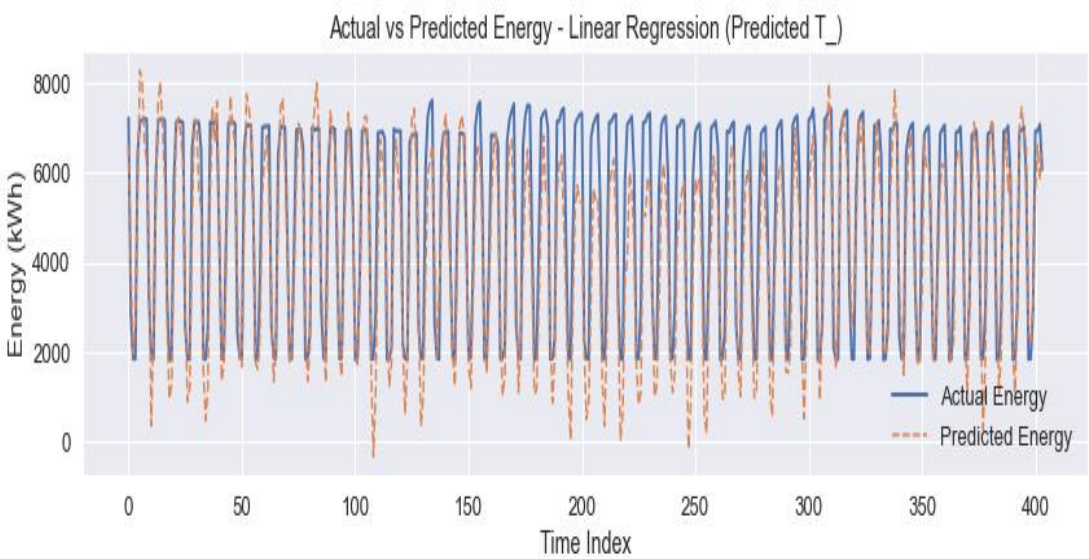

Figure 5A: Time-Series Comparison

Received: August 09, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

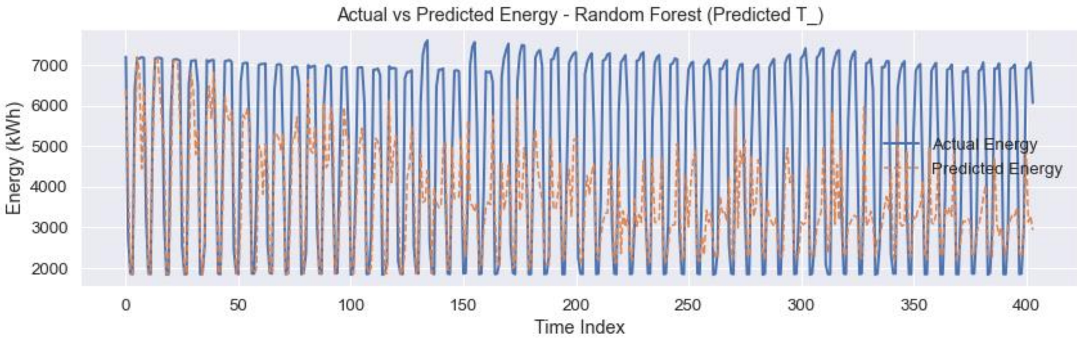

Figure 5B: Time-Series Comparison

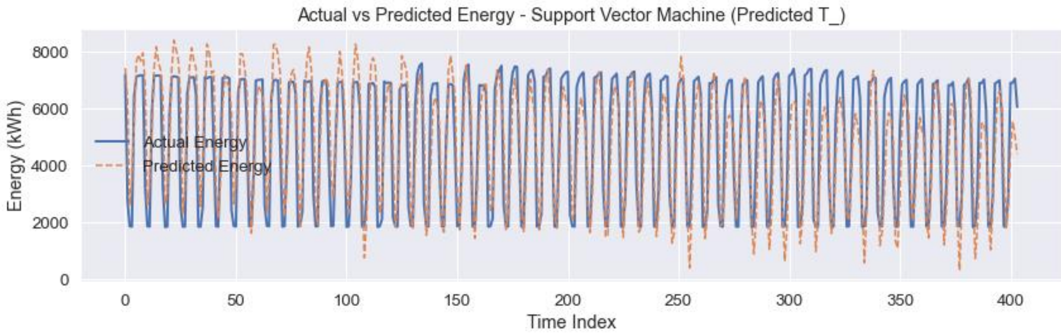

Figure 5C: Time-Series Comparison

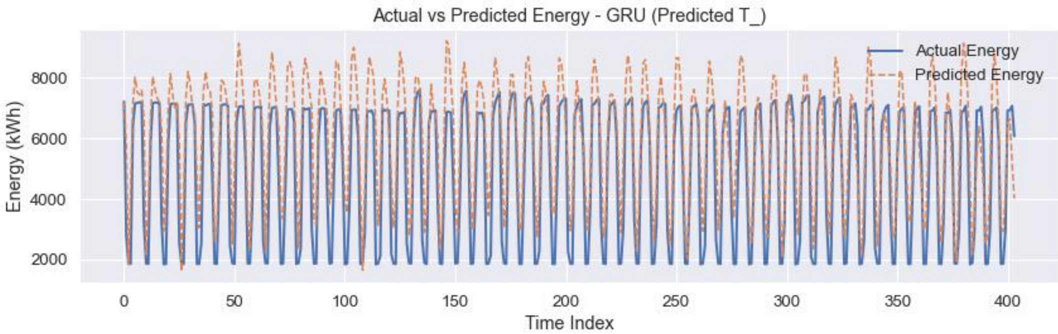

Figure 5D: Time-Series Comparison

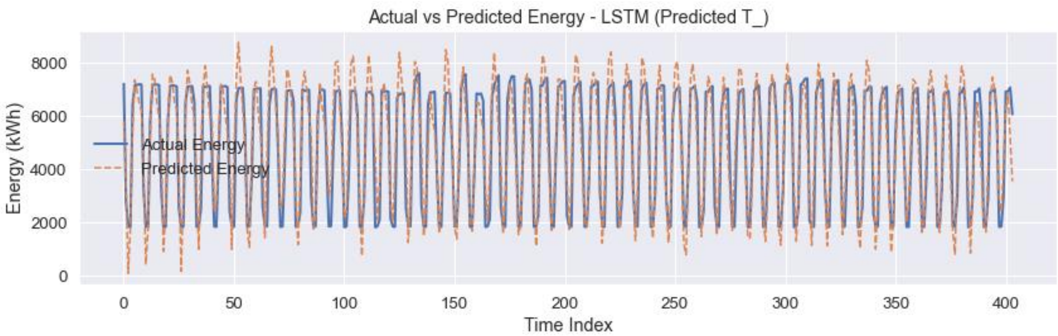

Figure 5E: Time-Series Comparison

Received: August 09, 2025

International Journal of Applied Mathematics
Volume 38 No. 4s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

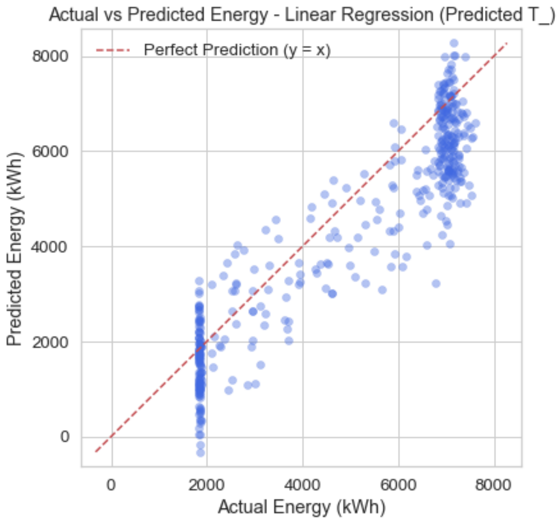

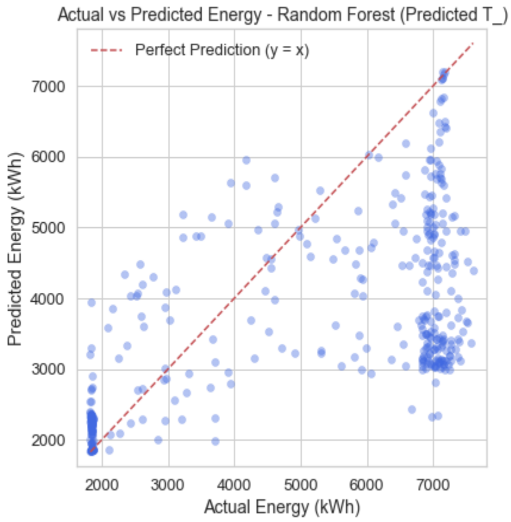

Figure 6A: Scatter Plot Comparison

Figure 6B: Scatter Plot Comparison

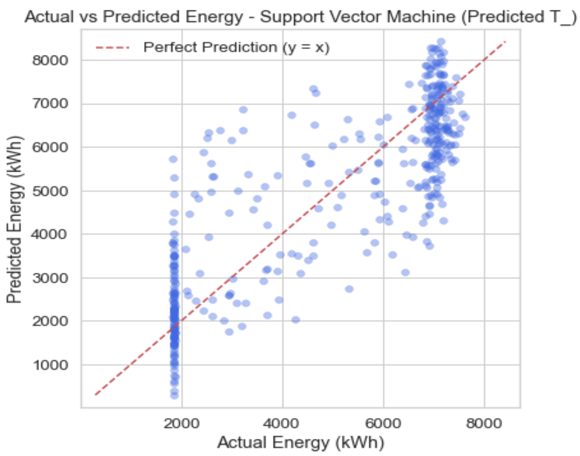

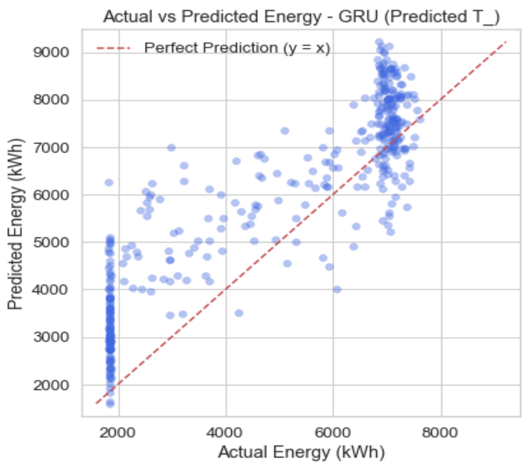

Figure 6C: Scatter Plot Comparison

Figure 6D: Scatter Plot Comparison

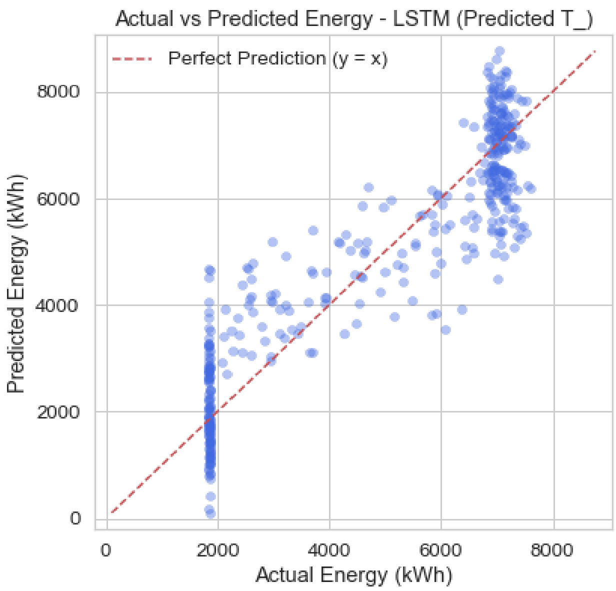

Figure 6E: Scatter Plot Comparison

Received: August 09, 2025

Volume 38 No. 4s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

These figures collectively emphasize that predictive accuracy (as measured by RMSE and R^2) and operational effectiveness (as reflected in energy savings) are not always aligned. Some models prioritize statistical fit, while others inadvertently optimize toward total consumption reduction. Therefore, selecting the appropriate model for a specific HVAC control or forecasting application should consider both predictive performance and energy impact.

In summary, LSTM and Linear Regression appear to offer the most balanced performance, with LSTM showing resilience to predicted inputs and LR offering simplicity and consistency. Random Forest may be better suited for energy-centric applications where accuracy is secondary to cost or energy savings.

6. Discussion

The results of this study reveal important insights into the applicability, robustness, and energy-saving potential of machine learning (ML) and deep learning (DL) models in HVAC forecasting under realistic constraints. By evaluating model performance across two input conditions—actual and predicted indoor temperatures—we provide a nuanced understanding of model behavior in both ideal and practical scenarios.

First and foremost, models utilizing actual temperature data consistently achieved higher accuracy as measured by R^2 and lower RMSE values. This finding aligns with existing literature that confirms the advantage of complete and accurate input information in supervised learning contexts. For instance, Linear Regression and Random Forest models both achieved R^2 values above 0.88 when fed with actual temperatures, while GRU and LSTM also performed well with R^2 scores near 0.87 and 0.86, respectively. These results validate the predictive power of ML and DL models in a best-case scenario where sensor inputs are available in real-time.

However, the predicted temperature scenario—arguably more reflective of real-world building control applications—yielded more complex outcomes. While LSTM maintained strong performance ($R^2 = 0.821$) even with predicted inputs, its ability to translate accuracy into energy savings was limited. Conversely, the Random Forest model, which exhibited poor R^2 (0.077), paradoxically achieved the highest energy savings (26.77%). This discrepancy suggests that statistical accuracy alone may not fully explain energy efficiency outcomes, and alternative metrics such as cost savings or cumulative energy reduction may be more relevant in operational contexts.

One possible explanation for these contrasting results is the nature of the learning objectives and optimization criteria within each model. DL models like LSTM and GRU are generally optimized for minimizing error across the entire dataset, which may lead to conservative predictions that stay close to mean values. While this improves metrics like RMSE and R^2, it may fail to identify aggressive energy-saving opportunities. On the other hand, Random Forests can overfit certain data patterns, resulting in predictions that align more with energy-saving targets even if they deviate from observed trends. This divergence illustrates a crucial trade-off between fidelity and impact in building energy forecasting.

Another noteworthy observation is the performance consistency of Linear Regression across both scenarios. While it lacked the sophistication of more complex models, it consistently produced stable results with low error and positive energy savings. This reinforces the argument that simpler models may offer valuable performance-to-complexity trade-offs, especially in smaller buildings or systems with limited computational capacity.

The findings also raise practical considerations for smart building system designers. In environments where real-time temperature measurements are unavailable or costly, reliance on forecasted temperature inputs is inevitable. Our results show that even in such conditions, LSTM models can maintain acceptable predictive performance. However, integrating these models into real-time control systems will require attention to robustness, retraining schedules, and sensor drift.

Lastly, visual analysis through scatter plots and time-series comparisons confirmed the variability in prediction behavior among models. LSTM and GRU maintained close alignment with actual energy values but showed lag in response to sharp load changes. Random Forest captured some of these fluctuations better, possibly contributing to its higher energy savings despite lower overall accuracy.

In summary, this discussion underscores the importance of evaluating both statistical performance and domain-specific outcomes like energy savings. It advocates for a holistic model assessment strategy that incorporates accuracy, interpretability, computational cost, and real-world feasibility—an essential perspective for transitioning HVAC prediction models from research to deployment.

7. Conclusion

This study evaluates a range of machine learning (ML) and deep learning (DL) models for HVAC energy forecasting using both predicted and actual temperature inputs. By simulating real-world scenarios where future temperature values may be unavailable, we implement a two-stage architecture using a high-resolution dataset from Oak Ridge National Laboratory. Five models—Linear Regression, Random Forest, Support Vector Machine, GRU, and LSTM—were assessed based on prediction accuracy and energy savings.

Results show that models using actual temperature data deliver superior accuracy (R^2 up to 0.884), while models using predicted temperatures still provide competitive performance. Notably, LSTM maintained high accuracy even under input uncertainty, making it a strong candidate for real-time applications. Random Forest achieved the highest energy savings (26.77%) despite a low R^2, highlighting the importance of evaluating models on both statistical and practical performance.

These findings suggest that DL models like LSTM and GRU are well-suited for predictive control environments, while simpler models such as Linear Regression offer stable, interpretable performance. The divergence between accuracy and energy efficiency underscores the need to include domain-specific metrics when assessing model suitability.

Future work may incorporate external data sources, such as weather forecasts, and explore integration with control strategies like model predictive control (MPC). Overall, the proposed two-stage framework offers a robust, scalable solution for intelligent HVAC forecasting and supports the development of energy-efficient, data-driven building management systems.

References

[1] M. Jayashankara, P. Shah, A. Sharma, P. Chanak, and S. K. Singh, "A Novel Approach for Short-Term Energy Forecasting in Smart Buildings," *IEEE Sens J*, vol. 23, no. 5, pp. 5307–5314, Mar. 2023, doi: 10.1109/JSEN.2023.3237876. [CrossRef] [Google Scholar] [Publisher Link]

[2] D. Arias-Requejo, B. Pulido, M. M. Keane, and C. J. Alonso-Gonzalez, "Clustering and deep-learning for energy consumption forecast in smart buildings," *IEEE Access*, vol. 11, pp. 128061–128080, 2023, doi: 10.1109/ACCESS.2023.3331329. [CrossRef] [Google Scholar] [Publisher Link]

[3] R. Mathumitha, P. Rathika, and K. Manimala, "Intelligent deep learning techniques for energy consumption forecasting in smart buildings: a review," *Artif Intell Rev*, vol. 57, no. 2, Feb. 2024, doi: 10.1007/S10462-023-10660-8. [CrossRef] [Google Scholar] [Publisher Link]

[4] M. Cai, … S. R.-2018 I., and undefined 2018, "Impact of HVAC set point adjustment on energy savings and peak load reductions in buildings," *ieeexplore.ieee.orgM Cai, S Ramdaspalli, M Pipattanasomporn, S Rahman, A Malekpour, SR Kothandaraman2018 IEEE International Smart Cities Conference (ISC2), 2018•ieeexplore.ieee.org*, Accessed: Jul. 23, 2025. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/8656738/

[5] F. P. S. Almeida, M. Castelli, and N. Côrte-Real, "Leveraging Feature Sets and Machine Learning for Enhanced Energy Load Prediction: A Comparative Analysis," *Emerging science journal*, vol. 8, no. 6, pp. 2120–2143, Dec. 2024, doi: 10.28991/ESJ-2024-08-06-01. [CrossRef] [Google Scholar] [Publisher Link]

[6] M. Jayashankara, P. Shah, A. Sharma, P. Chanak, and S. K. Singh, "A Novel Approach for Short-Term Energy Forecasting in Smart Buildings," *IEEE Sens J*, vol. 23, no. 5, pp. 5307–5314, Mar. 2023, doi: 10.1109/JSEN.2023.3237876. [CrossRef] [Google Scholar] [Publisher Link]

[7] S. Jiang and H. Hui, "Energy Consumption Probabilistic Prediction of HVAC Systems in Public Buildings Based on Deep Learning Fusion Model," *2024 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids, SmartGridComm 2024*, pp. 71–76, Sep. 2024, doi: 10.1109/SMARTGRIDCOMM60555.2024.10738104. [CrossRef] [Google Scholar] [Publisher Link]

[8] M. Gholamzadehmir, C. Del Pero, … S. B.-S. C. and, and undefined 2020, "Adaptive-predictive control strategy for HVAC systems in smart buildings–A review," *ElsevierM Gholamzadehmir, C Del Pero, S Buffa, R FedrizziSustainable Cities and Society, 2020•Elsevier*, Accessed: Jul. 23, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2210670720307009. [CrossRef] [Google Scholar] [Publisher Link]

[9] H. Gong, R. E. Alden, A. Patrick, and D. M. Ionel, "Forecast of Community Total Electric Load and HVAC Component Disaggregation through a New LSTM-Based

Method," *Energies (Basel)*, vol. 15, no. 9, pp. 2974–2974, Apr. 2022, doi: 10.3390/EN15092974. [CrossRef] [Google Scholar] [Publisher Link]

[10] Y. Chen, Y. Ye, J. Liu, L. Zhang, W. Li, and S. Mohtaram, "Machine Learning Approach to Predict Building Thermal Load Considering Feature Variable Dimensions: An Office Building Case Study," *Buildings*, vol. 13, no. 2, pp. 312–312, Jan. 2023, doi: 10.3390/BUILDINGS13020312. [CrossRef] [Google Scholar] [Publisher Link]

[11] M. Sakib, T. Siddiqui, S. Mustajab, R. M. Alotaibi, N. M. Alshareef, and M. Z. Khan, "An ensemble deep learning framework for energy demand forecasting using genetic algorithm-based feature selection," *PLoS One*, vol. 20, no. 1, pp. e0310465–e0310465, Jan. 2025, doi: 10.1371/JOURNAL.PONE.0310465. [Google Scholar] [Publisher Link]

[12] G. SHIRSAT, A. MUKHERJEE, and A. K. SHARMA, "Strategic load scheduling in smart home: Leveraging IOT and optimization algorithms for energy cost reduction," *Sigma Journal of Engineering and Natural Sciences*, 2024, doi: 10.14744/SIGMA.2024.00132. [CrossRef]

[13] N. Pateriya, N. Thakre, P. K. Kewat, and A. Vishwakarma, "IoT in the Energy Sector: Revolutionizing Energy Management and Consumption," *International Journal of Innovative Research in Computer and Communication Engineering*, vol. 11, no. 10, pp. 11344–11351, Oct. 2023, doi: 10.15680/IJIRCCE.2023.1110050. [CrossRef] [Publisher Link]

[14] F. P. S. Almeida, M. Castelli, and N. Côrte-Real, "Leveraging Feature Sets and Machine Learning for Enhanced Energy Load Prediction: A Comparative Analysis," *Emerging science journal*, vol. 8, no. 6, pp. 2120–2143, Dec. 2024, doi: 10.28991/ESJ-2024-08-06-01.

[15] H. Gong, R. E. Alden, A. Patrick, and D. M. Ionel, "Forecast of Community Total Electric Load and HVAC Component Disaggregation through a New LSTM-Based Method," *Energies (Basel)*, vol. 15, no. 9, pp. 2974–2974, Apr. 2022, doi: 10.3390/EN15092974.

[16] S. Mahjoub, L. Chrifi-Alaoui, B. Marhic, and L. Delahoche, "Predicting Energy Consumption Using LSTM, Multi-Layer GRU and Drop-GRU Neural Networks," *Sensors*, vol. 22, no. 11, pp. 4062–4062, May 2022, doi: 10.3390/S22114062. [CrossRef] [Google Scholar] [Publisher Link]

[17] P. V. B. Ramos, S. M. Villela, W. N. Silva, and B. H. Dias, "Residential energy consumption forecasting using deep learning models," *Appl Energy*, vol. 350, Nov. 2023, doi: 10.1016/J.APENERGY.2023.121705. [CrossRef] [Google Scholar] [Publisher Link]

[18] Y. Chen, Y. Ye, J. Liu, L. Zhang, W. Li, and S. Mohtaram, "Machine Learning Approach to Predict Building Thermal Load Considering Feature Variable Dimensions: An Office Building Case Study," *Buildings*, vol. 13, no. 2, pp. 312–312, Jan. 2023, doi: 10.3390/BUILDINGS13020312.

[19] Y. Yoon, S. Jung, P. Im, and A. Gehl, "Datasets of a Multizone Office Building under Different HVAC System Operation Scenarios," *Sci Data*, vol. 9, no. 1, pp. 1–11, Dec. 2022, doi: 10.1038/S41597-022-01858-

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

6;SUBJMETA=166,2790,4077,639;KWRD=ENERGY+MODELLING,ENGINEERING. [CrossRef] [Google Scholar] [Publisher Link]

[20] N. Luo, Z. Wang, D. Blum, C. Weyandt, N. B.-S. data, and undefined 2022, "A three-year dataset supporting research on building energy management and occupancy analytics," *nature.comN Luo, Z Wang, D Blum, C Weyandt, N Bourassa, MA Piette, T HongScientific data, 2022•nature.com*, Accessed: Jul. 23, 2025. [Online]. Available: https://www.nature.com/articles/s41597-022-01257-x. [CrossRef] [Google Scholar] [Publisher Link]

[21] D. Zhao *et al.*, "Data-driven online energy management framework for HVAC systems: An experimental study," *Appl Energy*, vol. 352, p. 121921, Dec. 2023, doi: 10.1016/J.APENERGY.2023.121921. [CrossRef] [Google Scholar] [Publisher Link]

[22] EIA (US Energy Information Administration). Monthly Energy Review, EIA (2021).

Received: August 09, 2025