EPH-International Journal of Applied Science

Volume - 12, Issue - 03
July-2026
DOI-https://doi.org/10.69980/pe563v71

DATA-DRIVEN PREDICTION OF BUILDING ENERGY LOADS

Meghna Modi

M.Tech student, artificial intelligence and data science , Department: of computer science , RGPV, Indore
Pin code: 452002, Email Id: meghnamodi3102002@gmail.com

Received:- 01-06-2026, Revised:- 02-07-2026, Accepted:- 09-07-2026, Published:- 16-07-2026

Abstract
In the modern building world, energy-load prediction is a vital aspect of energy optimisation, demand management and sustainable energy use. This study developed and compared interpretable machine-learning models for predicting hourly building energy consumption across two buildings with different operational characteristics. Multi-year hourly observations were preprocessed through temporal feature engineering, incorporating historical energy consumption, meteorological conditions, and calendar-related variables. The four supervised regression models, namely Linear Regression, Random Forest, Histogram Gradient Boosting, and Extreme Gradient Boosting (XGBoost), were evaluated using an 80:20 chronological train-test split. Predictive performance was assessed using Mean Absolute Error (MAE), Root Mean Square Error (RMSE), Mean Absolute Percentage Error (MAPE), and the coefficient of determination (R^2), while SHapley Additive exPlanations (SHAP) were applied to interpret feature importance. Histogram Gradient Boosting consistently achieved the highest prediction accuracy, obtaining an MAE of 8.063, RMSE of 14.151, MAPE of 5.131%, and R^2 of 0.9661 for Building A, and an MAE of 3.574, RMSE of 5.481, MAPE of 10.303%, and R^2 of 0.9752 for Building B. SHAP analysis identified one-hour, daily, and weekly historical energy consumption as the dominant predictors, whereas meteorological variables contributed comparatively less. The results presented indicate that temporal feature engineering and gradient-based ensemble learning can provide an accurate, interpretable and efficient framework for short-term building energy-load forecasting and intelligent energy management.

Keywords: *Building energy prediction, machine learning, gradient boosting, SHAP analysis, load forecasting*

Introduction

Buildings are a significant part of the world's energy consumption and are one of the largest sources of resource use and GHGs. Due to rapid urbanization, the growth of building volume and the demand for indoor comfort, all residential, commercial and institutional buildings are seeing their energy needs grow, and efficient energy management is a key factor for sustainable development (Berardi, 2017). As the urbanization of the world continues to rise in the next decades, building energy consumption will surely increase even more, underscoring the importance of implementing intelligent and effective strategies taking both energy efficiency and environmental sustainability into account (Güneralp et al., 2017). Energy efficiency, technology and operational management are also expected to be the most important factors in meeting future energy demand from buildings, and scenario analyses over the long-term are consistent with this (Levesque et al., 2018). Therefore, increasingly accurate information about the energy use of buildings is important for energy planning and decision-making. Due to the interplay between weather and occupancy behaviour, operational schedules, equipment characteristics, and building design, building energy demand is highly dynamic. In-depth studies on energy end uses and the factors that influence them have shown that knowledge of these interactions is essential for the development of an effective conservation strategy and better building performance (González-Torres et al., 2022). At the urban scale, optimisation of building layouts and operational characteristics has also been recognised as an effective approach for improving energy balance while maintaining environmental quality (Natanian et al., 2019). It is also revealed from recent studies that the optimisation of energy efficiency in residential buildings can lead to substantial environmental and economic gains due to its effect on energy savings and reduction of energy-related emissions (Abdunnabi et al., 2023).

Over the last few years, the spread of new sensor technologies, smart meters and digital monitoring systems has shifted the focus of building energy analysis from a statistical approach to modelling using data. Historical consumption data has been combined with environmental and operational data with advances in the field of data science, and predictive models have been able to capture the complex relationships that are not easily represented using classical analytical methods (Molina-Solana et al., 2017). Machine-learning methods are therefore an essential part of contemporary building energy management systems, as they can easily learn from observed data the nonlinear patterns. A wide variety of studies have tested the effectiveness of machine-learning algorithms in estimating the energy use in buildings. Most of the early reviews showed that, in general, data-driven methods are superior to physics-based methods, especially for short-term applications in forecasting (Seyedzadeh et al., 2018). The same result has been seen in the general review of building energy forecasting by using machine learning models, which have shown high adaptability to the different operational conditions with fewer assumptions on physical building parameters (Amasyali & El-Gohary, 2018). Ensemble learning algorithms and boosting techniques have also been found to be more accurate predictors than traditional regression models, especially when the model has to reflect the nonlinear relationships between the predictor variables (Bourdeau et al., 2019).

With the recent advances in deep learning, more opportunities have been opened for the energy forecasting field, which can automatically extract complex temporal patterns from a huge amount of data. Long short-term memory networks, recurrent neural networks and other deep learning models have been shown to be successful for modeling sequential energy consumption data, especially if the data are available over a long period of time (Runge & Zmeureanu, 2021). However, even with these improvements, it is still possible to choose a suitable prediction model based on data types, computational performance, and the need to be interpretable. Building energy modelling is also a topic of great interest at an urban scale since cities are interested in optimizing the energy performance of their buildings and in enhancing their energy resilience. Along with high-quality algorithms, high-quality input data, temporal resolution, weather variables, and weather operational information used in the modelling process have been identified as being of utmost importance for reliable prediction (Hong et al., 2020). In recent years, data-driven models incorporating energy consumption history and environmental data have shown significant gains in prediction accuracy, highlighting the importance of temporal feature engineering and environmental information in the

field of practical prediction (Mishra et al., 2024).
Previous research has shown that machine-learning techniques can be used to create energy prediction models that are effective, but there are a number of limitations. It can be difficult to compare models; many times the models are based on a single predictive algorithm or the models are tested with various preprocessing procedures. Temporal feature engineering is often applied in an ad hoc manner, and the effect of the historical energy variables, weather conditions and calendar effects is not always tested in a common modelling framework. Moreover, few studies have considered several regression algorithms in the same context with the same chronological validation methods and have included the interpretability of the models to determine which variables are the most important in terms of prediction accuracy. Overcoming these limitations is crucial for the development of prediction frameworks that are not only accurate, but also transparent and applicable in the context of practical building energy management.
The present study aims to develop a data-driven framework for predicting hourly building energy loads from multi-year observations of energy demands and the meteorological and calendar-related variables. Temporal feature engineering is used to include the short-term and medium-term consumption trends, and then assess the performance of Linear Regression, Random Forest, Histogram Gradient Boosting and Extreme Gradient Boosting on the same chronological validation approach. The performance of the model is evaluated using the standard regression parameters, and SHapley Additive exPlanations (SHAP) are used to measure the contribution of each predictor and enhance the model interpretability. The overall goal of the study is to develop an accurate and explainable model that can facilitate intelligent building energy management and operational decision-making through this integrated approach.

2. Methodology
2.1 Research Design
The research methodology used in this study was a quantitative and data-driven research design to build and test machine-learning models to predict the hourly building energy load. To achieve methodological consistency and the objective comparison of the performances, two independent building datasets were analysed through the same modelling workflow. The workflow involved in the model building comprised the preprocessing of the data, temporal feature engineering, feature selection, chronological data partitioning, model building, predictive evaluation, and model interpretation. In order to maintain the temporal order of observations and represent realistic forecasting situations, a chronological validation strategy was developed to minimize information leakage and enhance the model's evaluation. The general methodology was developed to find a predictive model for building energy consumption that is accurate and interpretable.

2.2 Data Source
Data for this study were provided by an open repository by Mariano (2024). The dataset includes multi-year hourly records of building energy consumption and meteorological data and calendar information for energy forecasting. These variables are adequate to explore the correlation between the operating parameters and the energy demand of the building and to create predictive models based on the data (Mariano, 2024).

2.3 Data Preprocessing
Raw data were first analyzed to determine data completeness, consistency in time, and equal data sampling intervals. There are several variables that correspond to energy consumption, the weather data and calendar data were synchronized with the energy data prior to creating the predictor variables. A number of descriptive statistics were computed to describe the energy and environmental variable distributions. To represent recurrent operational patterns, time-based attributes like hour of day, day of week, month, weekend status and holiday indicators were encoded. To capture the short-term and medium-term temporal dependency, lagged versions of energy variables were generated for energy history over one hour, one day and one week, as well as rolling mean and rolling standard deviation of energy variables over 24-hour and 168-hour periods, respectively. To correct for the fact that the first seven days were not complete, observations were

lagged.

2.4 Feature Selection and Chronological Data Partitioning
To remove redundant information and retain informative predictors feature selection was carried out prior to model development. All constants that didn't change throughout the training data were excluded, and the remaining predictive variables were screened by correlation in order to reduce the level of multicollinearity. Twenty-one and 19 predictors were chosen for Building A and Building B, respectively, according to their predictive relevance. The datasets were then split into a training set (80% of observations) and a testing set (20% of observations) in a chronological split. This temporal segregation maintained the natural appearance of the data in the model and prevented the use of future observations for model training.

2.5 Model Development
Four supervised regression models were implemented for comparison: Linear Regression, Random Forest, Histogram Gradient Boosting and Extreme Gradient Boosting (XGBoost). The same predictor variables and chronological training dataset were used for each model, so they are equally comparable. The trained models were used to predict hourly energy loads for the independent testing period, allowing the predictive capability to be assessed consistently between buildings. The computational effort spent in training and inference for the model was also measured, and they were also compared with predictive performance.

2.6 Performance Evaluation
The methods used for the evaluation of the models' performance were Mean Absolute Error (MAE), Root Mean Square Error (RMSE), Mean Absolute Percentage Error (MAPE), and coefficient of determination (R^2). The lower the MAE, RMSE, and MAPE values, the better the prediction accuracy, and the higher the R^2 values, the more agreement there was between the observed and predicted energy consumption. The complementary evaluation measures allowed for a comprehensive comparison of the accuracy and robustness of the models in the buildings data set and helped to identify the best prediction model.

2.7 Model Interpretation
Following the evaluation of the model's predictions, SHapley Additive exPlanations (SHAP) analysis was conducted to enhance the transparency of the model. The mean absolute SHAP values for individual features were calculated to quantify the impact of the individual features in model predictions and for ranking feature importance. This interpretation step allowed the identification of the most influential time, operational and environmental parameters affecting building energy consumption, which, together with the predictive evaluation, provided an interpretation of the model behaviour.

3. Results
3.1 Dataset Characteristics
After removing data with missing target values and dealing with inconsistencies in time, the cleaned datasets comprised 43,742 hourly observations for Building A and 43,778 hourly observations for Building B. The two datasets included building energy consumption as well as meteorological and calendar variables and were both based on the same time periods (January 2016 until December 2020). The energy demand for buildings was compared, with the results of Building A having significantly higher energy demands than Building B, with mean hourly consumptions of 165.59 and 66.64 units, respectively. The distribution of seasonal weather variables was similar between the two buildings, suggesting that the energy demand differences were mainly caused by operational characteristics of the buildings and not by the climatic conditions. Table 1 shows the descriptive statistics of the cleaned datasets.

Table 1. Descriptive Statistics of Building Energy Consumption and Environmental Variables

Building	Variable	Observations	Mean	Standard Deviation	Minimum	Median	Maximum
Building A	Energy Consumption	43,742	165.585	72.328	47.933	131.882	469.756
Building A	Heating Degree Days	43,742	6.374	5.536	0.000	6.205	19.690
Building A	Cooling Degree Days (0°C Base)	43,742	13.001	6.996	0.000	12.095	29.010
Building A	Cooling Degree Days (10°C Base)	43,742	4.597	5.302	0.000	2.095	19.010
Building A	Precipitation	43,742	1.365	3.066	0.000	0.080	46.170
Building A	Relative Humidity	43,742	65.740	17.290	24.090	65.880	100.000
Building A	Mean Air Temperature	43,742	12.500	7.263	−2.620	11.700	28.630
Building A	Minimum Air Temperature	43,742	6.852	5.878	−5.860	6.410	20.960
Building A	Maximum Air Temperature	43,742	19.147	8.466	2.760	18.130	38.660
Building A	All-Sky Solar Radiation	43,742	4.552	2.352	0.270	4.440	8.670

Building A	Holiday Indicator	43,742	0.403	0.490	0.000	0.000	1.000
Building B	Energy Consumption	43,778	66.636	45.190	5.647	49.776	259.150
Building B	Heating Degree Days	43,778	6.367	5.535	0.000	6.190	19.690
Building B	Cooling Degree Days (0°C Base)	43,778	13.010	6.997	0.000	12.110	29.010
Building B	Cooling Degree Days (10°C Base)	43,778	4.603	5.304	0.000	2.110	19.010
Building B	Precipitation	43,778	1.368	3.067	0.000	0.080	46.170
Building B	Relative Humidity	43,778	65.740	17.287	24.090	65.890	100.000
Building B	Mean Air Temperature	43,778	12.509	7.263	−2.620	11.710	28.630
Building B	Minimum Air Temperature	43,778	6.860	5.878	−5.860	6.430	20.960
Building B	Maximum Air Temperature	43,778	19.155	8.467	2.760	18.170	38.660
Building B	All-Sky Solar Radiation	43,778	4.554	2.352	0.270	4.450	8.670
Building B	Holiday Indicator	43,778	0.402	0.490	0.000	0.000	1.000

3.2 Comparative Performance of Prediction Models

An identical chronological training-testing scheme was used to assess 4 regression algorithms. The prediction errors were lowest and the coefficient of determination was highest for Histogram

Gradient Boosting in both buildings. For Building A, the model attained an RMSE of 14.151, MAE of 8.063, MAPE of 5.13%, and an R^2 of 0.9661. The same superiority was seen with Building B with the values being 5.481, 3.574, 10.30%, and 0.9752, respectively. XGBoost generated pretty close results, but not as accurate as HGB, while RF and LR had comparatively higher prediction errors. In Table 2, the quantitative comparison of all prediction models evaluated is presented. Table 2. There is currently no information to compare the performance of the machine-learning models.

Table 2. Comparative Predictive Performance of Machine-Learning Models

Building	Model	Training Observations	Testing Observations	Features	MAE	RMSE	MAPE (%)	R^2	Training Time (s)	Best Model
Building A	Histogram Gradient Boosting	34,859	8,715	21	8.063	14.151	5.131	0.9661	2.178	Yes
Building A	XGBoost	34,859	8,715	21	10.357	17.176	7.293	0.9500	2.896	No
Building A	Random Forest	34,859	8,715	21	12.956	18.555	10.105	0.9416	35.096	No
Building A	Linear Regression	34,859	8,715	21	13.316	21.241	8.775	0.9235	0.337	No
Building B	Histogram Gradient Boosting	34,888	8,722	19	3.574	5.481	10.303	0.9752	2.150	Yes
Building B	XGBoost	34,888	8,722	19	4.459	6.295	13.872	0.9673	2.696	No
Building B	Random Forest	34,888	8,722	19	6.276	8.046	24.944	0.9466	35.451	No
Building B	Linear Regression	34,888	8,722	19	8.850	11.217	30.143	0.8962	0.080	No

Figure 1A shows that Histogram Gradient Boosting achieved the lowest RMSE among all competing models for Building A. Figure 1B shows that the same model also provided the smallest prediction error for Building B, demonstrating stable predictive performance across different building profiles.

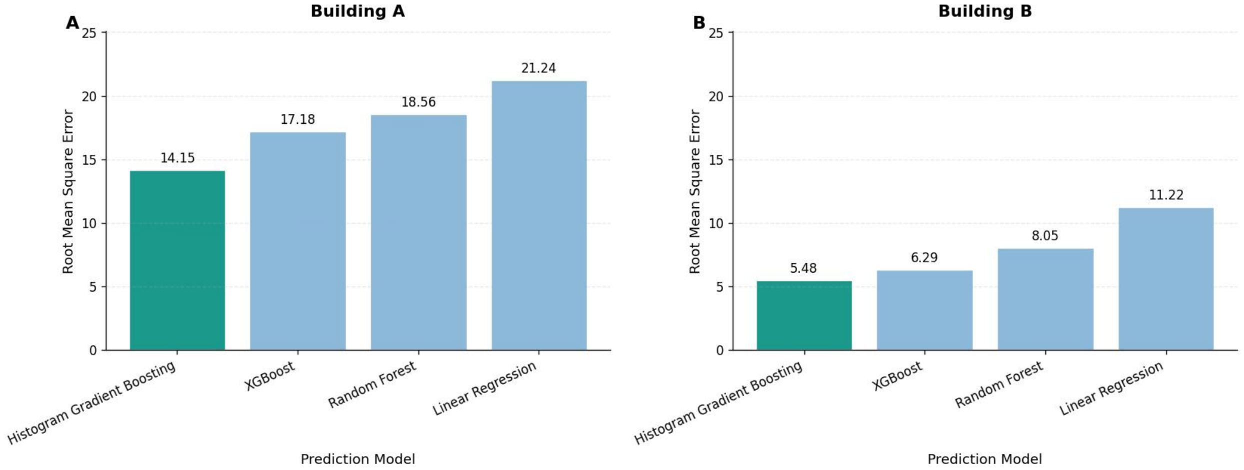

Figure 1. Comparative Prediction Performance of Machine-Learning Models. (A) Root mean square error comparison for Building A. (B) Root mean square error comparison for Building B.

3.3 Prediction Accuracy and Model Explainability

The results of the predictions on the hourly energy consumption were obtained by a visual examination that showed good agreement between the computed and observed consumption in the entire experimental period. The Histogram Gradient Boosting (HGB) model was able to capture both regular consumption cycles and periods of high consumption with relatively low deviations during periods of sudden change. The model was found to be a good capture of temporal variation in building energy demand in that the two curves are very close. The observed and predicted hourly energy loads for Building A for the testing period are shown in Figure 2A. Similarly, the prediction accuracy for Building B is presented in Figure 2B; the predicted energy loads were well aligned with the observed ones.

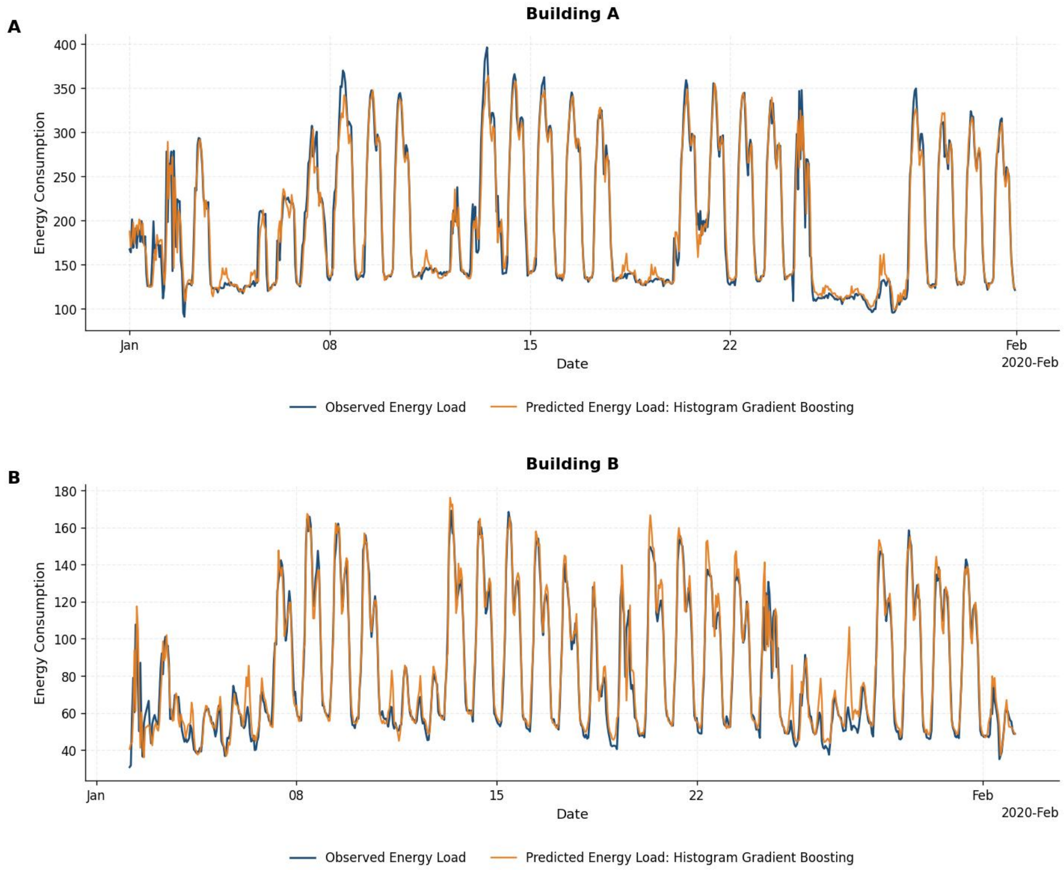

Figure 2. Observed and Predicted Hourly Building Energy Loads. (A) Observed and predicted energy consumption for Building A using Histogram Gradient Boosting. (B) Observed and predicted energy consumption for Building B using Histogram Gradient Boosting.

The XGBoost model was used to perform SHAP analysis to make the prediction models easier to interpret. Previous-hour energy use, as well as daily and weekly lagged energy uses, cyclical hourly energy uses, and calendar energy indicators, were identified as the most important predictors for both buildings. The other weather variables (such as temperature and heating degree days) were also related to the prediction accuracy, but their explanatory effect was smaller than historical energy variables. The results show that the main factor affecting the energy use of buildings was the short-term temporal dependence in both datasets. The 10 most influential predictors, ranked by the mean absolute SHAP value, are listed in Table 3.

Table 3. SHAP-Based Importance of the Leading Predictors

Building	Rank	Predictor	Mean Absolute SHAP Value
Building A	1	Energy Consumption Lag 1 Hour	49.209
Building A	2	Energy Consumption Lag 24 Hours	8.743
Building A	3	Energy Consumption Lag 168 Hours	6.851
Building A	4	Hour Cosine	6.707
Building A	5	Hour Sine	4.412
Building A	6	Day of Week Sine	2.080
Building A	7	Energy Consumption Rolling Mean 168 Hours	1.489
Building A	8	Holiday Indicator	1.444
Building A	9	Day of Week Cosine	1.439
Building A	10	Energy Consumption Rolling Mean 24 Hours	1.255
Building B	1	Energy Consumption Lag 1 Hour	30.752
Building B	2	Energy Consumption Lag 24 Hours	4.920
Building B	3	Energy Consumption Lag 168 Hours	3.024
Building B	4	Hour Sine	2.823
Building B	5	Hour Cosine	2.302
Building B	6	Energy Consumption Rolling Mean 24 Hours	1.640
Building B	7	Holiday Indicator	1.398
Building B	8	Day of Week Cosine	0.820
Building B	9	Day of Week Sine	0.807
Building B	10	Energy Consumption Rolling Mean 168 Hours	0.657

The previous-hour energy consumption was the most important predictor for Building A, while the daily and weekly lag variables and hourly cyclic components were the next most important variables, as can be seen in Figure 3A. The historical energy demand and the temporal operating schedules were also more important for predicting energy load in Building B, as seen in Figure 3B.

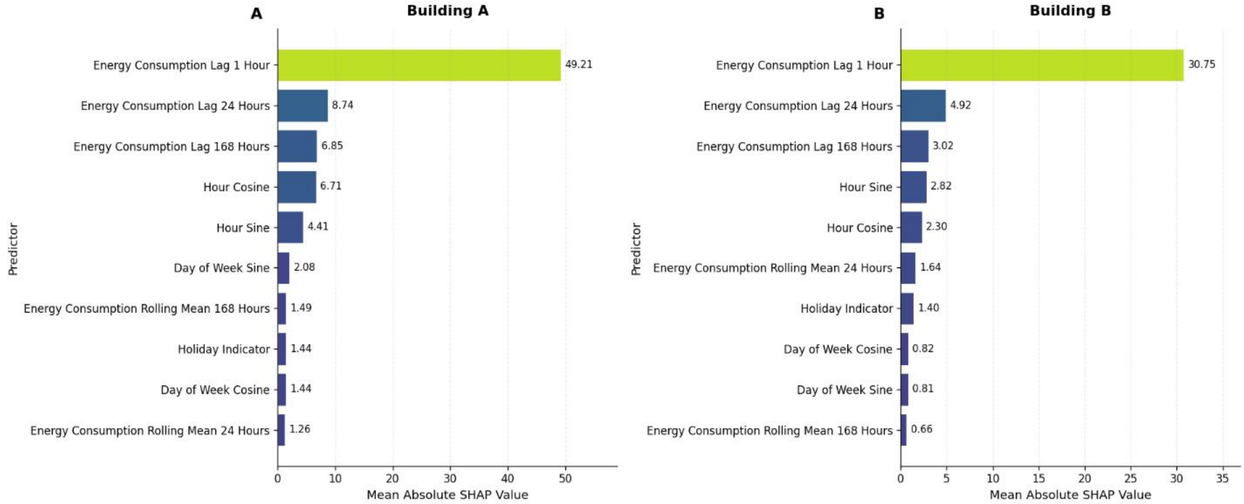

Figure 3. SHAP-Based Importance of Energy-Load Predictors. (A) Mean absolute SHAP values for the ten most influential predictors in Building A. (B) Mean absolute SHAP values for the ten most influential predictors in Building B.

4. Discussion

They show that energy loads of a building can be well predicted from the historical loads and the temporal attributes and environmental conditions when analysed in a unified machine-learning framework. Histogram Gradient Boosting shows the best results for both buildings; R^2 values for Buildings A and B were 0.9661 and 0.9752, respectively. The lower values of MAE and RMSE suggest that the model was able to capture the nonlinear relationships better than the other models, specifically Linear Regression, Random Forest and XGBoost. The average ranking was obtained over two buildings having very different consumption levels, which further validates the use of gradient-based ensemble learning for short-term building energy forecasting. In addition, the model needed significantly less training time compared to the Random Forest model, providing a good balance between prediction accuracy and computational cost. Absolute energy demand alone did not prove to determine the difficulty of prediction, as differences between the two buildings indicate. The absolute errors were higher for Building A, as it had a much higher mean hourly consumption with a larger variation. Building B had lower values of MAE and RMSE, but a higher MAPE value due to the fact that percentage errors are more significant when the observed values are relatively low. Regardless of these differences in demand profiles, both models had good R^2 values, indicating that the models accounted for a large proportion of the temporal variation. Good fit between the curves further suggests that the framework chosen captured the daily cycles, weekly patterns, low load periods and most high load events. SHAP analysis gave more support to the idea that the most important information for prediction was from recent consumption history. The previous hour's energy demand had the highest value in both buildings, with 24-hour and 168-hour lags following. A good overall ranking is gained because of good persistence and the daily and weekly reoccurrence of building operations. Other sine and cosine variables were also important, and they also supported the point that energy consumption was on a regular intraday rhythm. Holiday, weekend and day-of-week indicators were used to reflect changes that occurred due to lower or different occupancy. The meteorological variables contributed less than the historical load features, indicating that rather than dominating it, operational continuity was more significant in the short-term prediction, and weather conditions were used to enhance the predictions.

The significant influence of historical data via monitoring is consistent with previous research showing that temporal consumption patterns that are well-approximated by measured building data can enable accurate prediction of load with monitoring (Zheng et al., 2017). The superiority of

Histogram Gradient Boosting over Linear Regression also lends credibility to a previous report that carefully designed nonlinear machine learning models are more effective in capturing complex relationships between energy demand and its predictors than traditional linear models (Seyedzadeh et al., 2019). Advanced nonlinear models have been proven to deliver significant accuracy improvements through comparative studies on heating-load prediction, but the accuracy of these models has been found to be related to the configuration of the algorithms and the type of data available (Le et al., 2019). The current results then add to that conclusion by demonstrating that a relatively simple model with boosting, which is more efficient than linear models or tree-based models, is able to outperform them without being very complex, while using an energy model that is only hourly. Deep sequential models have also been demonstrated to accurately predict building electricity consumption while enabling operational optimisation (Zhou et al., 2022), which captures the long-term dependencies. This is the current result that shows that a large portion of this temporal information can be captured in an interpretable tabular modelling framework, with the help of engineered lag and rolling features.

The successful application of chronological validation is consistent with the city-based forecasting research where the temporal ordering is essential to the assessment of realistic future prediction, instead of just random interpolation of past observations (Wang et al., 2021). Multi-scale energy modelling frameworks that use engineering knowledge and machine learning to enhance operational and planning decisions (Nutkiewicz et al., 2018) are also supported by the results. This agreement implies that the modelling process could be employed in a more extended building or city energy management system as a practical tool. The research works on combining heterogeneous information sources has shown that the quality of energy estimation is enhanced by incorporating environmental and operational information (Pan & Zhang, 2020). The SHAP results complement this evidence, as they reveal that not all information sources are equally important; the historic consumption was the most explanatory, while the latter ones were of secondary importance: temperature, degree days, humidity and solar radiation. Machine learning has also been used to make predictions on building performance and retrofit assessment, with recent applications demonstrating the utility of data-driven prediction for identifying potential energy savings and implementing energy measures to achieve these goals (Ali et al., 2024). Further work on AI forecasting suggests that precise forecasting of demand can enhance the sustainability of institutional and urban planning (Chowdhury, 2025).

The results can be applied in intelligent building operation. Accurate hourly predictions can help with demand scheduling, peak-load management, anomaly detection, and heating/cooling/electricity system coordination. The extensive influence of the one-hour and daily lags suggests that it is possible to derive valuable predictions from a smart meter's data, even when detailed building information is unavailable. These variables for the calendar and weather can then be added to make the responsiveness to operational changes more effective. The added value for Facility Managers is the interpretability of the model, in which the prediction is preceded by clear evidence of what is driving the demand. The SHAP rankings can be useful to identify the difference between routine temporal changes and weather-driven changes, and unusual consumption behaviour. The Histogram Gradient Boosting performed well in terms of accuracy and computational demands, allowing for frequent retraining within the typical energy-management framework.

There are a few restrictions which should be noted. Results are only applicable to the two buildings analysed and may not reflect the performance of all building types and/or operational conditions. Due to a lack of occupancy data and detailed equipment-level loads, energy consumption was only analysed as an aggregate target instead of individual loads for heating, cooling, lighting, and appliances. Since the best-performing model was the Histogram Gradient Boosting, the interpretation was done using XGBoost.

The system should be tested in more building types in the future, both in the residential, commercial and institutional sectors. Using occupancy, equipment schedules, indoor conditions and energy system variables can enhance predictions during sudden changes in operation. In addition, further improvements to reliability may be achieved by employing methods of rolling validation, probabilistic forecasting, and by directly explaining the selected model. Evaluating transfer learning and cross-building generalisation would also help to understand how well a model developed in one building can be adapted to a different one.

5. Conclusion

A data-driven framework that integrates historical energy consumption, temporal features, weather data and machine-learning techniques was developed and used to accurately predict hourly building energy loads by employing a unified approach. The best results were obtained by the Histogram Gradient Boosting model, with R^2 values of 0.9661 and 0.9752 for the two buildings, and with the lowest prediction errors compared to the other models such as Linear Regression, Random Forest and XGBoost. The results indicate that the complex energy-demand patterns can be effectively captured by the nonlinear ensemble models in the different buildings with different energy-demand patterns. The largest component of predictive information was the historical energy use. The one-hour, 24-hour and 168-hour lag variables were always ranked higher than all the meteorological variables, and the cyclical hourly and calendar variables were able to capture the recurring operational behaviour. SHAP analysis accounted for the added benefit of enhancing model transparency by providing insight into the role that recent demand and temporal schedules played in making predictions. The proposed framework provides a practical foundation for short-term load forecasting, demand management, anomaly detection, and intelligent building operation. Since it requires only smart-meter and weather data, which are readily available, it is feasible to deploy without detailed physical building models. More generalisation in different building types, occupancy and end-use would further confirm its generalisability and usefulness in operation.

References

1. Abdunnabi, M., Etiab, N., Nassar, Y. F., El-Khozondar, H. J., & Khargotra, R. (2023). Energy savings strategy for the residential sector in Libya and its impacts on the global environment and the national economy. Advances in Building Energy Research, 17(4), 379-411.
2. Ali, U., Bano, S., Shamsi, M. H., Sood, D., Hoare, C., Zuo, W., ... & O'Donnell, J. (2024). Urban building energy performance prediction and retrofit analysis using a data-driven machine learning approach. Energy and Buildings, 303, 113768.
3. Amasyali, K., & El-Gohary, N. M. (2018). A review of data-driven building energy consumption prediction studies. Renewable and Sustainable Energy Reviews, 81, 1192-1205.
4. Berardi, U. (2017). A cross-country comparison of the building energy consumptions and their trends. Resources, Conservation and Recycling, 123, 230-241.
5. Bourdeau, M., qiang Zhai, X., Nefzaoui, E., Guo, X., & Chatellier, P. (2019). Modeling and forecasting building energy consumption: A review of data-driven techniques. Sustainable Cities and Society, 48, 101533.
6. Chowdhury, B. R. (2025). AI-Powered Forecasting and Optimization of Energy Consumption in the USA: Machine Learning Approaches for Sustainable Urban and Institutional Development. Euro Vantage journals of Artificial intelligence, 2(2), 21-38.
7. González-Torres, M., Pérez-Lombard, L., Coronel, J. F., Maestre, I. R., & Yan, D. (2022). A review on buildings energy information: Trends, end-uses, fuels and drivers. Energy Reports, 8, 626-637.
8. Güneralp, B., Zhou, Y., Ürge-Vorsatz, D., Gupta, M., Yu, S., Patel, P. L., ... & Seto, K. C. (2017). Global scenarios of urban density and its impacts on building energy use through 2050. Proceedings of the National Academy of Sciences, 114(34), 8945-8950.
9. Hong, T., Chen, Y., Luo, X., Luo, N., & Lee, S. H. (2020). Ten questions on urban building energy modeling. Building and Environment, 168, 106508.
10. Le, L. T., Nguyen, H., Dou, J., & Zhou, J. (2019). A comparative study of PSO-ANN, GA-ANN, ICA-ANN, and ABC-ANN in estimating the heating load of buildings' energy efficiency for smart city planning. Applied Sciences, 9(13), 2630.
11. Levesque, A., Pietzcker, R. C., Baumstark, L., De Stercke, S., Grübler, A., & Luderer, G. (2018). How much energy will buildings consume in 2100? A global perspective within a scenario framework. Energy, 148, 514-527.
12. Mariano, D. (2024). Building energy consumption data (Version 2) [Data set]. Mendeley Data. https://doi.org/10.17632/mzkyh37mtr.2

13. Mishra, A., Lone, H. R., & Mishra, A. (2024). DECODE: Data-driven energy consumption prediction leveraging historical data and environmental factors in buildings. Energy and Buildings, 307, 113950.
14. Molina-Solana, M., Ros, M., Ruiz, M. D., Gómez-Romero, J., & Martín-Bautista, M. J. (2017). Data science for building energy management: A review. Renewable and Sustainable Energy Reviews, 70, 598-609.
15. Natanian, J., Aleksandrowicz, O., & Auer, T. (2019). A parametric approach to optimizing urban form, energy balance and environmental quality: The case of Mediterranean districts. Applied Energy, 254, 113637.
16. Nutkiewicz, A., Yang, Z., & Jain, R. K. (2018). Data-driven Urban Energy Simulation (DUE-S): A framework for integrating engineering simulation and machine learning methods in a multi-scale urban energy modeling workflow. Applied energy, 225, 1176-1189.
17. Pan, Y., & Zhang, L. (2020). Data-driven estimation of building energy consumption with multi-source heterogeneous data. Applied Energy, 268, 114965.
18. Runge, J., & Zmeureanu, R. (2021). A review of deep learning techniques for forecasting energy use in buildings. Energies, 14(3), 608.
19. Seyedzadeh, S., Rahimian, F. P., Glesk, I., & Roper, M. (2018). Machine learning for estimation of building energy consumption and performance: a review. Visualization in Engineering, 6(1), 5.
20. Seyedzadeh, S., Rahimian, F. P., Rastogi, P., & Glesk, I. (2019). Tuning machine learning models for prediction of building energy loads. Sustainable Cities and Society, 47, 101484.
21. Wang, Z., Hong, T., Li, H., & Piette, M. A. (2021). Predicting city-scale daily electricity consumption using data-driven models. Advances in Applied Energy, 2, 100025.
22. Zheng, Z., Zhuang, Z., Lian, Z., & Yu, Y. (2017). Study on building energy load prediction based on monitoring data. Procedia Engineering, 205, 716-723.
23. Zhou, X., Lin, W., Kumar, R., Cui, P., & Ma, Z. (2022). A data-driven strategy using long short term memory models and reinforcement learning to predict building electricity consumption. Applied Energy, 306, 118078.