Energy Savings in Buildings Using Predictive Analysis

Leland Sion
University of Victoria
lelandsion@gmail.com

Cole Westendorf
University of Victoria
ckwestendorf@gmail.com

***Abstract*—Effective energy management in buildings is essential for reducing operational costs, enhancing efficiency, and minimizing environmental impact. This paper explores the integration of machine learning techniques, specifically Long Short-Term Memory (LSTM) networks, to predict energy consumption patterns and optimize usage. By leveraging predictive energy modeling, buildings can reduce peak demand, shift nonessential loads, and enhance overall energy efficiency. The study examines the potential benefits of LSTM-based forecasting in enabling data-driven decision-making, leading to smarter and more sustainable energy management strategies.**

I. Introduction

Energy management in buildings is crucial for reducing costs, improving efficiency, and minimizing environmental impact. With the integration of machine learning models like Long Short-Term Memory (LSTM) networks, buildings can predict their energy consumption patterns and optimize usage accordingly. This paper explores how predictive energy modeling can help reduce peak demand, shift nonessential loads, and improve overall efficiency. To achieve this, we develop an LSTM-based forecasting model trained on historical energy data, incorporating key variables such as HVAC usage, occupancy trends, and environmental conditions. The study evaluates how these predictions enable demand response strategies, such as load shifting, battery discharge timing, and automated energy optimization. The paper first discusses the challenges of peak demand and energy forecasting, followed by an exploration of temperature prediction for HVAC efficiency. The methodology, results, and implications of predictive energy management are then analyzed, demonstrating how machine learning enhances energy savings in smart buildings.

A. Motivation

Peak demand refers to periods of highest energy consumption in a building, typically occurring when multiple systems—such as HVAC, lighting, and appliances—operate simultaneously. These peaks lead to higher utility costs due to demand charges, reduced HVAC efficiency from operating at full capacity, increased grid strain, and limitations in battery and renewable energy supply during peak loads. By accurately predicting peak demand, building managers can implement strategies to reduce consumption, shift loads to off-peak hours, and optimize HVAC performance. Leveraging a Long Short-Term Memory (LSTM) model trained on historical energy data, buildings can forecast energy demand for the next day and take preventive measures to minimize unnecessary energy use during peak hours. Effective strategies include load shifting, which reschedules nonessential activities (e.g., dishwashing, laundry, EV charging) to off-peak times, and pre-cooling and thermal storage, which cool buildings in advance to reduce HVAC load when occupancy is highest. Additional methods such as lighting optimization (adjusting brightness based on occupancy), battery discharge timing (strategic use of stored energy), and AI-driven automation further enhance efficiency. These approaches lower demand charges, improve sustainability, and optimize overall energy use. Forecasting HVAC demand allows for pre-conditioning spaces at optimal times, reducing sudden spikes and enhancing efficiency. Additionally, smart scheduling of energy-intensive operations (e.g., elevators, water heaters, commercial machinery) ensures they run during low-demand hours, reducing operational costs. Since utility providers often charge based on peak usage, predictive models help stagger high-energy processes, adjust HVAC settings, and manage non-essential loads, preventing unnecessary expenses and improving overall building energy management. Climate control, particularly heating and cooling, represents one of the largest energy expenditures in buildings. Temperature forecasting enables buildings to optimize HVAC operations efficiently. By analyzing temperature trends, predictive models facilitate pre-cooling or pre-heating strategies, allowing buildings to adjust HVAC operation in advance rather than reacting to external temperature fluctuations. This approach not only maintains occupant comfort but also reduces energy consumption.

B. Related Works

AI-Driven Energy Forecasting Using LSTM-Based Models: AI-based forecasting uses machine learning (ML) to predict a building's energy usage, allowing proactive adjustments to optimize consumption. Early approaches relied on physics-based simulations or classical ML models (e.g. regression, ARIMA, decision trees), which required expert-defined features and often struggled under volatile conditions like weather or occupancy changes [1]. Deep learning techniques have overcome many of these limitations. In particular, Long Short-Term Memory (LSTM) networks can automatically learn complex sequential patterns in energy data, making them highly effective for load forecasting [1]. Studies show that LSTM-based models consistently outperform traditional models –

for example, one achieved about 97% prediction accuracy, surpassing standard regression and decision-tree methods [1]. Such improvements in forecast precision are not just academic; they translate into better control. Modern Building Energy Management Systems (BEMS) leverage these accurate forecasts to make informed decisions (e.g. pre-cooling a building before occupancy or shifting loads), thereby maintaining comfort while minimizing waste [1]. In short, deep learning-driven forecasting has become a cornerstone of building energy optimization, enabling more efficient scheduling and resource allocation across HVAC, lighting, and other systems. Anomaly Detection Techniques for Energy Inefficiencies

Timely detection of anomalous energy behavior is equally vital for optimizing building performance. Anomalies – such as equipment malfunctions or user errors (e.g. an HVAC fault or lights left on) – appear as irregular usage patterns that, if uncorrected, lead to significant energy waste and even equipment damage [1]. AI-driven anomaly detection systems tackle this by learning normal consumption patterns and flagging deviations in real time. This field has evolved from simple rule-based or statistical thresholds to data-driven ML approaches that offer far more sensitivity and reliability [1]. In particular, deep learning models (using techniques like autoencoders or LSTM-based sequence models) excel at capturing the complex, non-linear relationships in building data, allowing them to identify unusual usage behaviors that traditional methods might miss [1]. One effective strategy is to combine forecasting with anomaly detection: for instance, a deep learning model can first filter out regular seasonal trends and use an LSTM to predict expected consumption, then flag any large discrepancy between the predicted and actual usage as an anomaly [2]. These AI-based systems provide reliable alerts to facility managers [1], so that faults or inefficiencies can be corrected quickly. By catching issues like a miscalibrated thermostat or a failing motor early, anomaly detection helps maintain optimal operations and prevents energy from being wasted needlessly. In practice, the Connect project leverages IoT infrastructure in commercial buildings to gather real-time energy data, which its LSTM-based AI engine uses for making short-term and long-term consumption forecasts. When the AI flags an unexpected surge or drop in usage (an anomaly), facility managers or automated controllers can be notified to take corrective action (e.g., investigate faulty equipment or adjust control strategies), closing the loop of smart building management.

C. Problem Definition

The reviewed literature also sheds light on several research gaps that Connect explicitly seeks to address. Himeur et al. point out enduring challenges in building energy anomaly detection, including the lack of (i) precise definitions of what constitutes an anomalous consumption event, (ii) annotated datasets for model training, (iii) unified metrics to evaluate detection performance, (iv) common platforms for reproducibility, and (v) measures for privacy preservation [3]. Connect tackles some of these gaps by adopting a clear operational definition of anomalies (e.g., significant deviation from the LSTM-predicted baseline for similar conditions) and by generating a repository of observed anomalies in its deployment building to serve as an evolving labeled dataset. In addition, Connect's evaluation framework combines forecast accuracy metrics (for the LSTM predictor) with anomaly detection precision/recall to provide a more unified assessment of energy management performance, aligning with calls for standardized metrics [**?**]. Another gap highlighted by Aguilar et al. is the need for developing autonomous cycles of data analysis tasks and better feature engineering in AI for smart buildings [2]. Currently, many solutions are fragmented, focusing on either prediction or control in isolation monitoring. Connect's architecture is designed to be more holistic: it blends real-time monitoring (via IoT), forecasting (via AI/LSTM), and a feedback mechanism for decision-making, thereby contributing to a more autonomous and integrated energy management loop. Moreover, the absence of techniques like online clustering for diagnostics in prior studies [2] suggests an opportunity for Connect's anomaly detection component to incorporate online learning, so it can adapt to new patterns (e.g., seasonal changes or shifts in building occupancy) without manual re-calibration. By addressing these research gaps – improved anomaly definitions, integrated analytics, and adaptability – the Connect project builds upon and extends the state of the art, as documented by the reviewed AI-in-building-energy research, to optimize energy consumption in commercial buildings.

II. Methodology

This study employs a data-driven approach to energy forecasting, leveraging deep learning techniques to predict building energy consumption. A Long Short-Term Memory (LSTM) neural network was selected due to its effectiveness in capturing long-term dependencies in time-series data. The model was trained on historical energy usage patterns, environmental factors, and occupancy trends to provide accurate predictions for key energy categories, including HVAC, Lighting, and Miscellaneous Electric Loads (MELS). The methodology consists of data preprocessing, feature selection, model architecture design, training, and performance evaluation. TensorFlow and Keras were used to implement the LSTM model, with optimized training parameters to ensure efficiency and accuracy in energy consumption forecasting.

AI Model and Training: This study employs a deep learning-based approach to predict building energy consumption using a Long Short-Term Memory (LSTM) neural network, implemented with TensorFlow and Keras. LSTMs, a specialized type of Recurrent Neural Network (RNN), are particularly well-suited for time-series forecasting due to their ability to capture long-term dependencies in sequential data. By analyzing historical energy usage patterns, the model provides forecasts for three key building energy categories: HVAC, Lighting, and Miscellaneous Electric Loads (MELS). The LSTM model is designed to process time-series energy data and predict future consumption patterns based on historical records, environmental conditions, and occupancy trends. The input features include historical energy consumption data segmented into HVAC, Lighting, and MELS categories, environmental conditions such as indoor and outdoor temperature, humidity, and other weather-related variables affecting heating and cooling demand, occupancy trends derived from office hours, human activity levels, and motion sensor data, as well as time-based features that capture the hour of the day, day of the week, and seasonal variations. The LSTM model architecture consists of multiple layers to extract temporal dependencies and refine predictions. The first LSTM layer contains 32 units with ReLU activation and return sequences enabled to pass information to subsequent layers. The second LSTM layer consists of 16 units with ReLU activation, capturing deeper sequential patterns in the data. The final dense output layer comprises three neurons corresponding to predicted energy consumption for HVAC, Lighting, and MELS. To ensure optimal model performance, the dataset undergoes several preprocessing steps, including data cleaning, where missing values are interpolated and outliers are filtered using statistical methods, and feature scaling using MinMaxScaler to normalize all input variables between 0 and 1, preventing bias in the learning process. Additionally, the dataset is transformed into sequences suitable for LSTM processing using TensorFlow's tf.data. Dataset API.

Model training is conducted over five epochs, a choice determined by validation loss trends. Training beyond five epochs resulted in increased memory consumption with minimal improvements in accuracy. Backpropagation Through Time (BPTT) is employed to optimize LSTM weights, and while early stopping was considered, it was not implemented due to the validation loss plateauing after five epochs. The TensorFlow ModelCheckpoint feature is used to store the best-performing model during training, ensuring robustness in deployment. The model is trained using the Mean Absolute Error (MAE) loss function, which is well-suited for energy forecasting, and the Adam optimizer, chosen for its balance of speed and stability. A batch size of 32 is used to maintain efficient computation without compromising learning stability. The dataset is split into 80 training data, consisting of one year's worth of historical records, and 20 testing data reserved for evaluation. Model performance is assessed using MAE scores for each energy category: HVAC MAE measures the accuracy in predicting heating and cooling demand, Lighting MAE evaluates the model's ability to forecast lighting energy use based on occupancy, and MELS MAE quantifies errors in predicting plug-load energy consumption. The ModelCheckpoint feature is also leveraged during training to save the best-performing model, ensuring consistency and reliability for deployment in real-world applications. The LSTM model was trained using 5 epochs with a batch size of 32. The decision to use 5 epochs was based on empirical testing, where further training beyond this point showed diminishing returns in reducing validation loss. During initial testing, training for more than 5 epochs led to hight use of ram and minimal gain to accuracy, the possibility to overfit was also considered, where the model performed well on training data but deteriorated on test data Early stopping was considered, but for efficiency, a fixed epoch number was chosen based on validation loss trends.

III. Results

To evaluate the model's accuracy in predicting hourly energy usage, the Mean Absolute Error (MAE) was used for both normalized and unnormalized data. The results are summarized below:

Graphical Representation of Model Performance To visualize the model's effectiveness, the following graphs illustrate actual vs. predicted energy usage for each category.

The graph shows that the model closely follows actual HVAC consumption trends. The effectiveness of temperature forecasting in HVAC optimization can be visualized through predictive models that estimate temperature fluctuations and dynamically adjust cooling or heating requirements. Anomaly Detection for Energy Waste Reduction Beyond forecasting, machine learning models detect anomalies in energy consumption, helping facility managers identify inefficiencies and prevent waste. Unexpected energy spikes may indicate equipment malfunctions, operational errors, or excessive consumption. Anomaly detection can trigger automated alerts for faulty HVAC or lighting systems, adjust systems in response to real-time inefficiencies, and flag abnormal patterns for proactive intervention. By identifying anomalies early, buildings can minimize energy waste, optimize performance, and reduce costs. The following visualization highlights energy usage anomalies, with red markers indicating potential inefficiencies or malfunctions.

Dataset Description The dataset used in this study is sourced from Dryad Digital Repository and was originally published in Nature Scientific Data [4]. The dataset spans 2020–2023, containing hourly energy consumption records for HVAC, lighting, and miscellaneous electrical loads (MELS) in a commercial buildings. It also includes environmental variables such as outdoor temperature, humidity, and wind speed, as well as occupancy data derived from motion sensors and scheduled building usage patterns. The dataset was preprocessed to ensure consistency, with missing values addressed using rolling mean interpolation and numerical features normalized between 0 and 1 using min-max scaling. The data was structured into sequences for time-series forecasting using LSTM networks, with an 80/20 train-test split applied, where 2020 pre-pandemic data was used for training to ensure stable energy consumption patterns. This dataset provides a comprehensive representation of real-world building energy use, supporting predictive modeling for energy efficiency and demand reduction. A detailed breakdown of dataset features and methodology is available in the Dryad Repository. To ensure data integrity and improve the performance of the predictive model, several preprocessing steps were applied before training. Missing values in critical variables, such as temperature readings and energy consumption records, were interpolated using statistical methods to maintain continuity in time-series data. Since energy usage data spans multiple

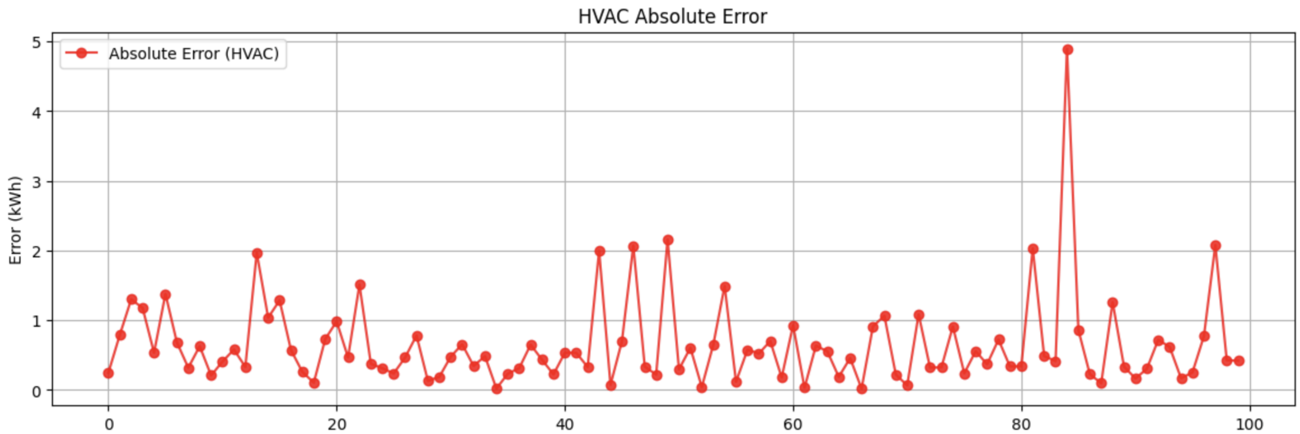

Fig. 1. HVAC Prediction: The HVAC system, which is one of the highest energy consumers, has an MAE of 0.63 kWh, meaning the model can predict energy demand with high accuracy. This enables pre-cooling or pre-heating strategies to be implemented efficiently.

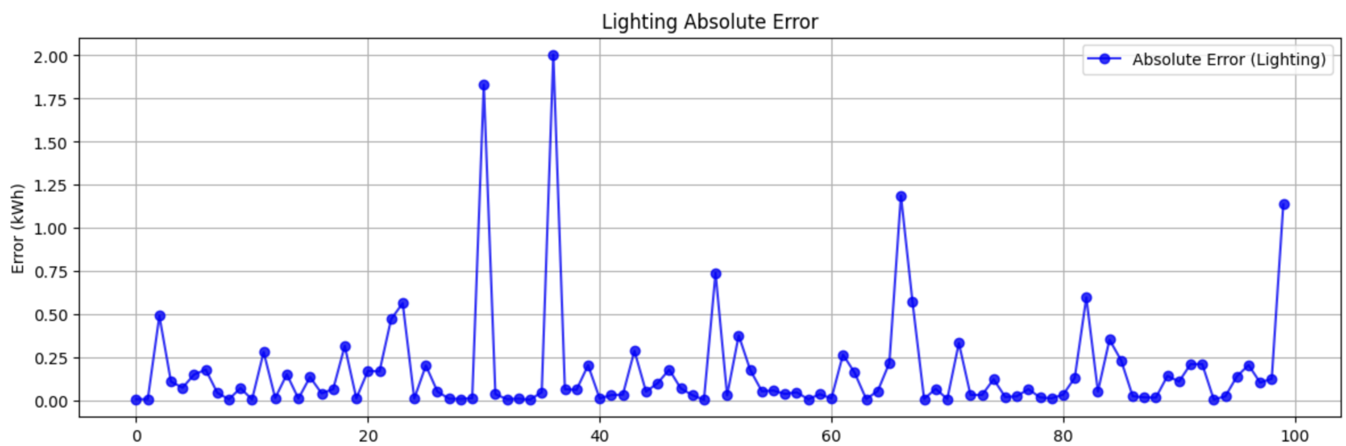

Fig. 2. Lighting Prediction: With an MAE of 0.11 kWh, the model accurately forecasts lighting needs, supporting smart dimming systems and occupancy-based adjustments.

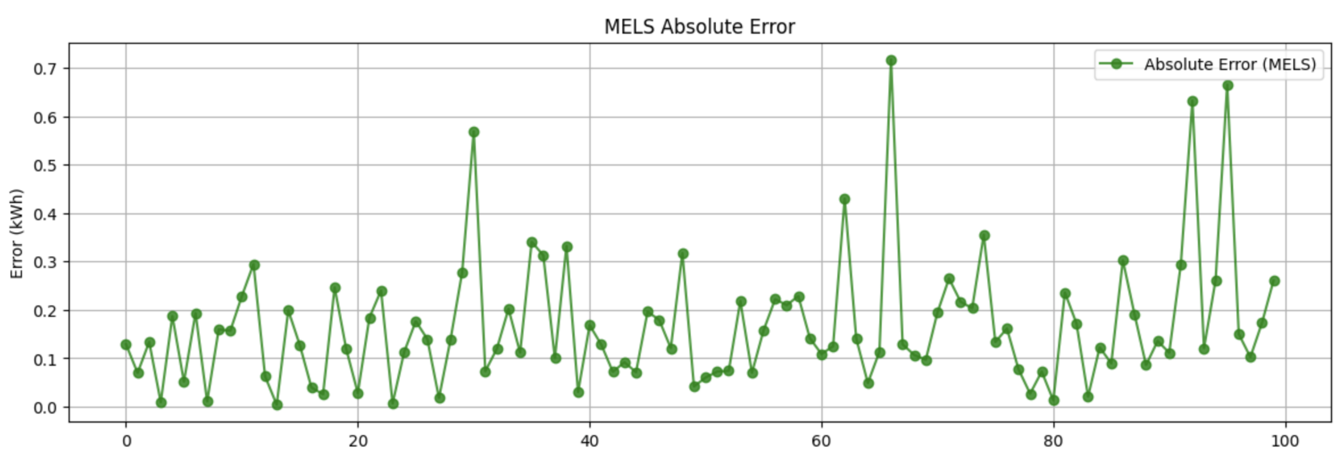

Fig. 3. MELS (Miscellaneous Electrical Loads) Prediction: The model achieves an MAE of 0.22 kWh, useful for detecting anomalies or optimizing device scheduling.

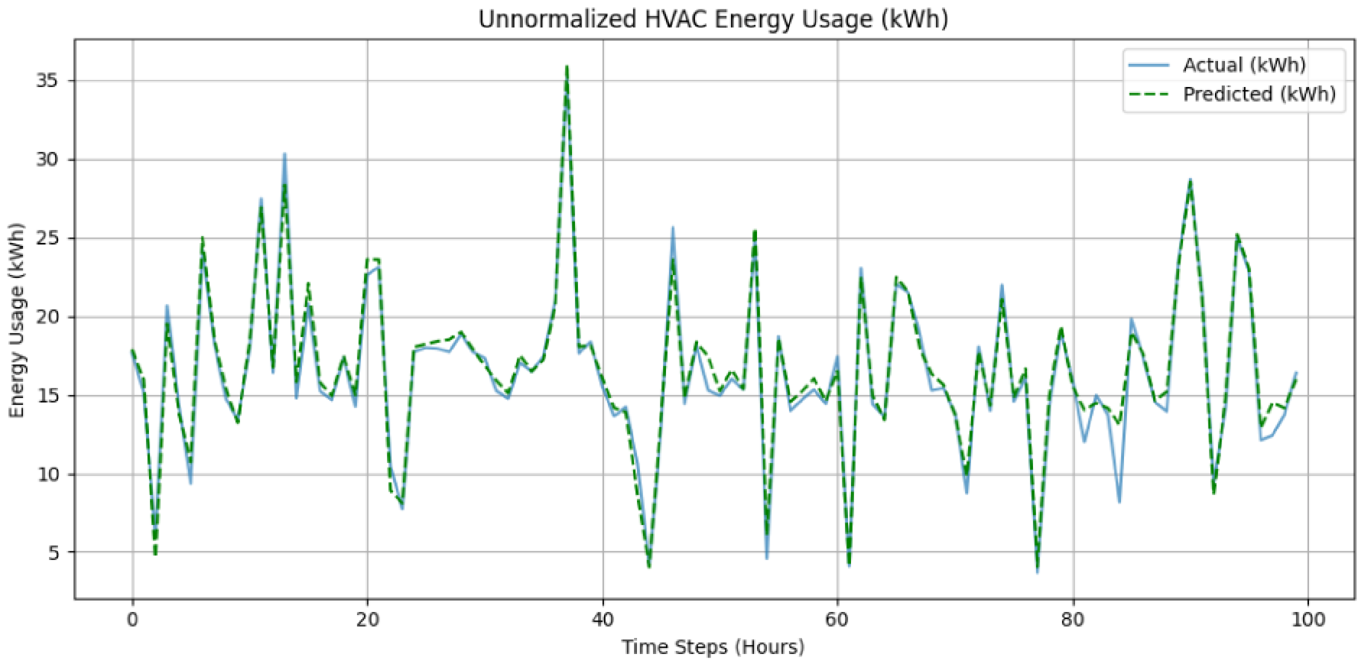

Fig. 4. Actual vs. predicted energy usage HVAC unnormalized.

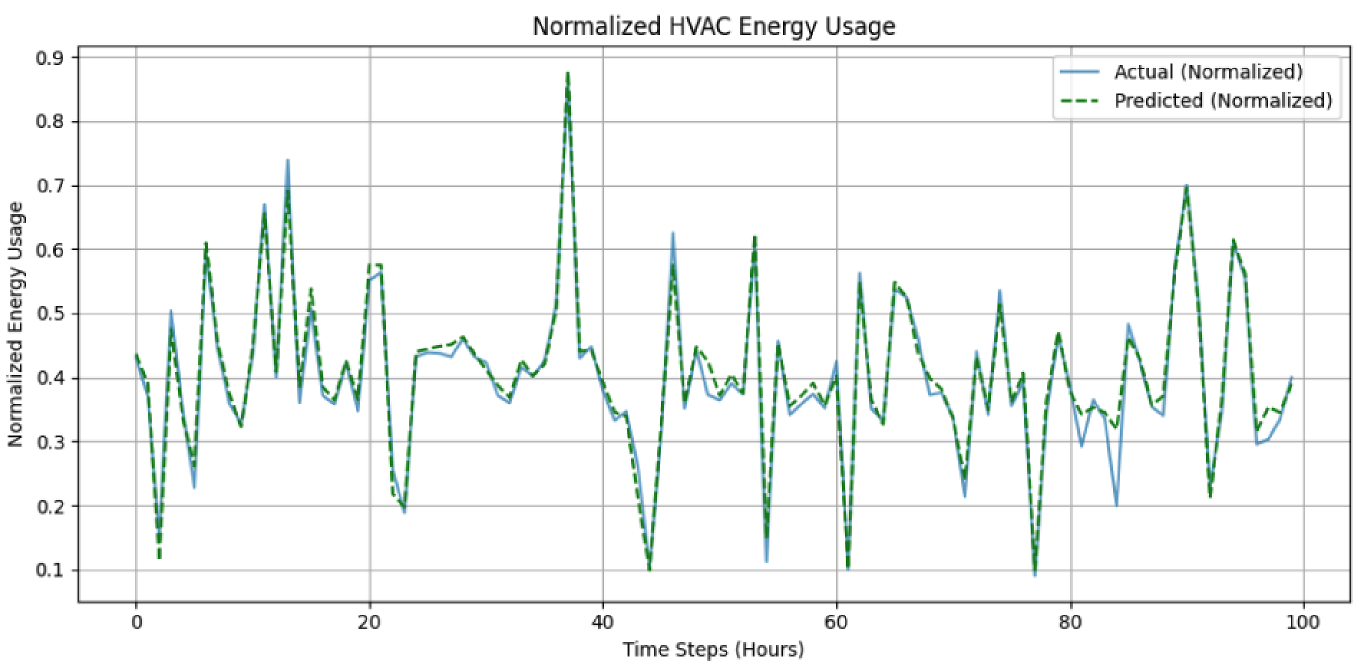

Fig. 5. Actual vs. predicted energy usage HVAC normalized.

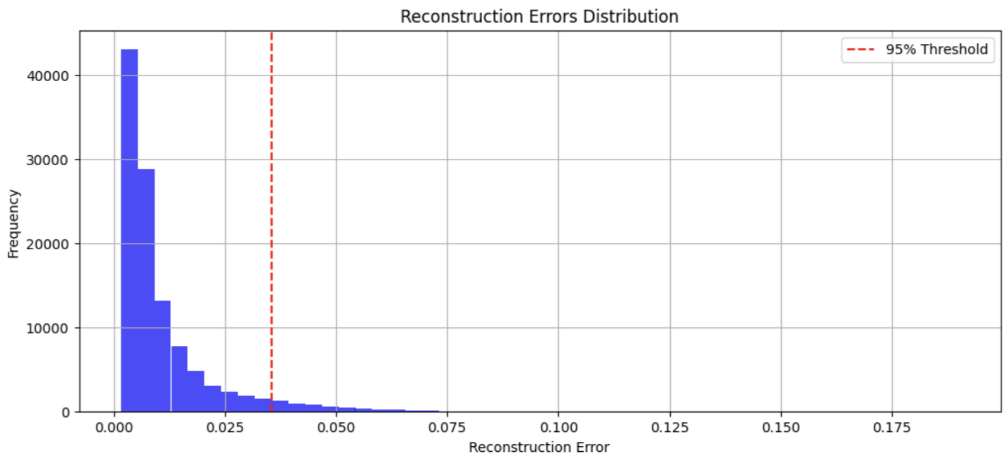

This histogram shows the distribution of reconstruction errors from the anomaly detection model. Most data points have low errors (indicating normal operation), but some outliers have high errors. The red dashed line represents the 90th percentile threshold, meaning that points above this line are flagged as anomalies.

```
Anomalies Detected: 2248 instances
```

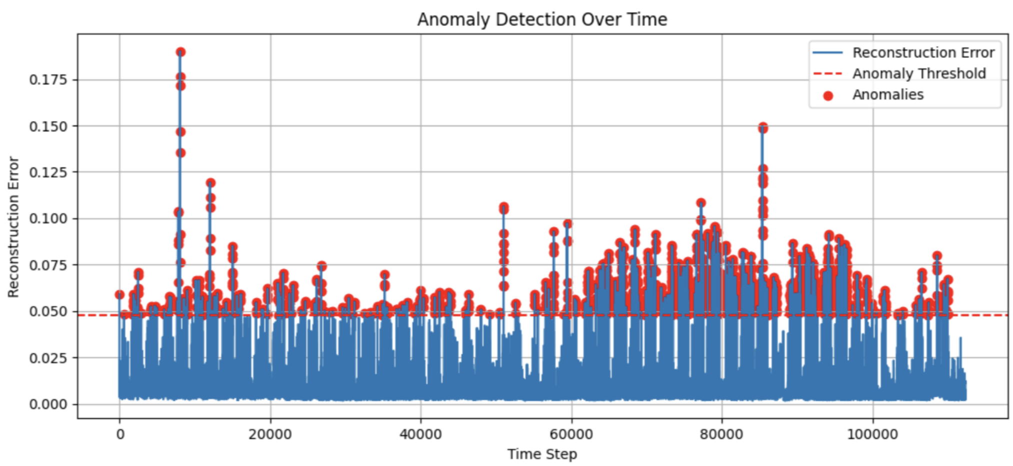

This time series graph represents the reconstruction errors over time. A higher error means the model struggled to reconstruct the input data, which could indicate unusual energy consumption patterns. The red dashed line represents the anomaly threshold, and the red dots mark timestamps where anomalies were detected.

magnitudes, feature scaling was performed using MinMaxScaler, normalizing all numerical values between 0 and 1 to prevent bias in model learning. Given the sequential nature of the dataset, the data was structured into time-series sequences suitable for Long Short-Term Memory (LSTM) networks using TensorFlow's tf.data.Dataset API, enabling the model to effectively capture temporal dependencies in energy consumption patterns. The dataset was split into training (80%) and testing (20%), with 2020 pre-pandemic data used for training to ensure that consumption patterns reflect a stable operational environment before disruptions introduced by occupancy and behavioral changes during COVID-19. The remaining 20% was reserved for testing and evaluation, allowing the model to generalize effectively to new data. By leveraging historical data and environmental conditions—including temperature, occupancy, and electrical loads—the model provides a holistic understanding of building energy usage, supporting the development of predictive optimization strategies for demand reduction and efficiency improvements.

The graphs above provide an overview of HVAC energy usage, variability, correlations, and daily patterns. The time-series plot (top left) highlights clustered HVAC demand shifts, while the daily energy and temperature variability graphs (top right) reveal fluctuations driven by environmental factors. The correlation heatmap (bottom left) shows strong dependencies between HVAC usage, temperature, and occupancy. Lastly, the hourly energy consumption plots(bottom right) illustrate stable HVAC demand and peak MELS usage in the evening. These insights help identify key drivers of energy consumption for predictive modeling.

To optimize energy consumption in buildings, a load-shifting strategy was implemented using predictive energy modeling. The model forecasts peak demand periods and dynamically reschedules non-essential loads (e.g., HVAC pre-cooling, lighting adjustments, and deferred appliance usage) to off-peak hours. This reduces energy costs by leveraging

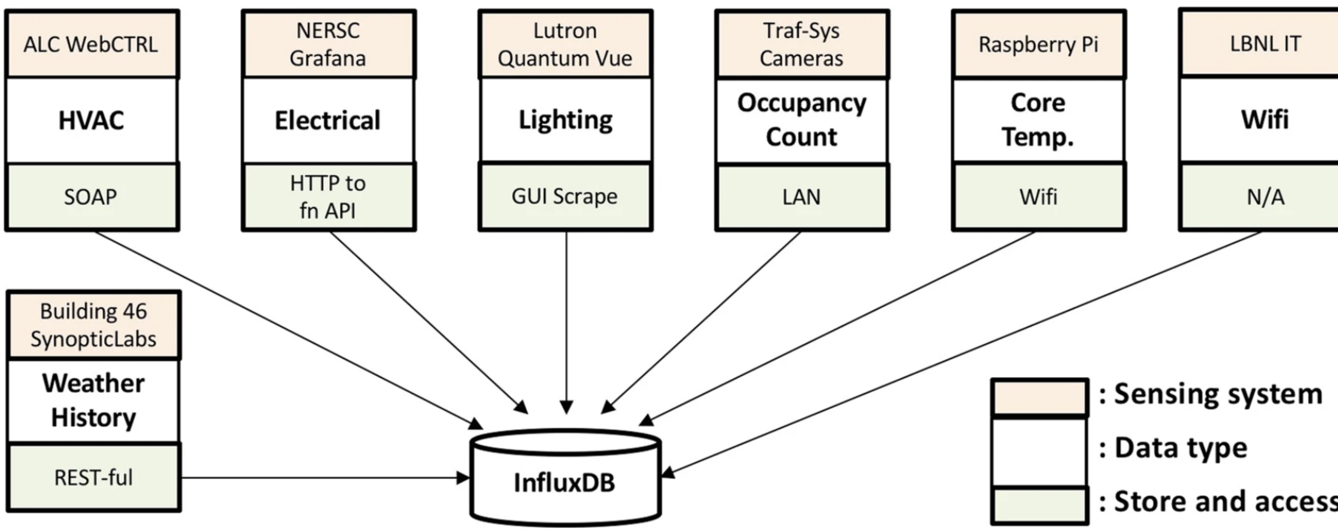

Fig. 7. The graph above indicate the structure of the data and how it was procured.

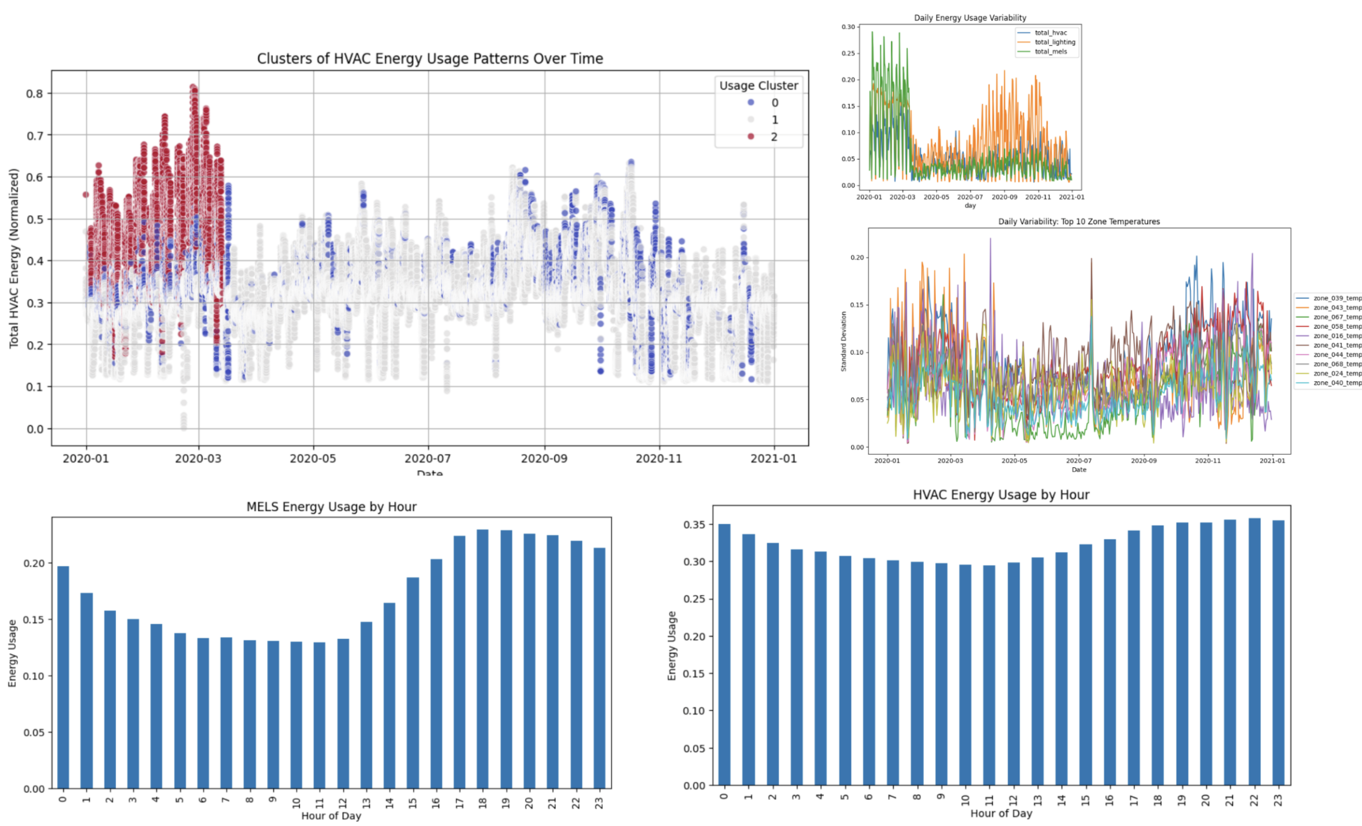

Fig. 8. Graphs showing energy usage clusters, variability and average energy usage from the building energy data.

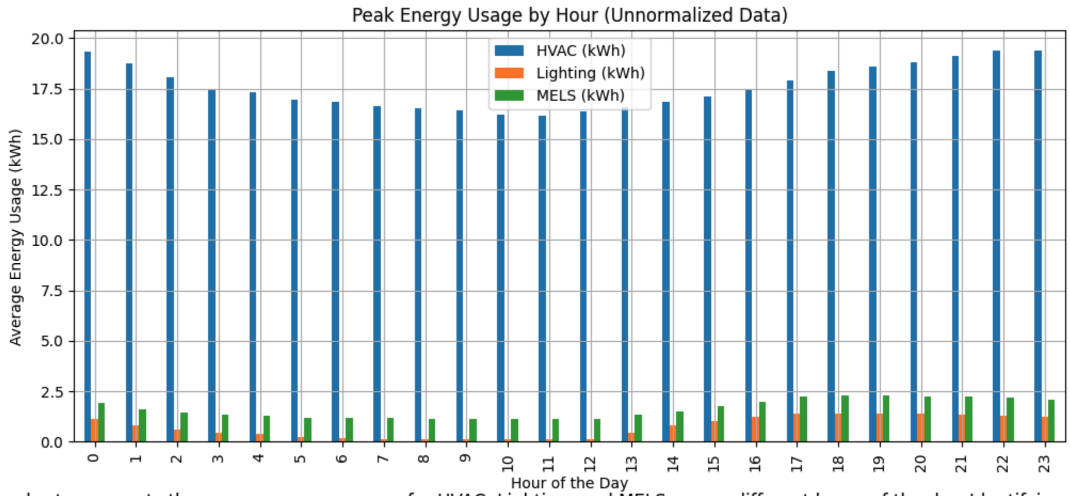

This bar chart represents the average energy usage for HVAC, Lighting, and MELS across different hours of the day. Identifying peak hours helps optimize load distribution and energy efficiency strategies.

```
Successfully extracted and unnormalized 24 hourly HVAC data points for 2020-01-15
Total Cost Before Load Shifting: 90.17
Total Cost After Load Shifting: 82.05
Total Savings: 8.12 (9.00% reduction)
```

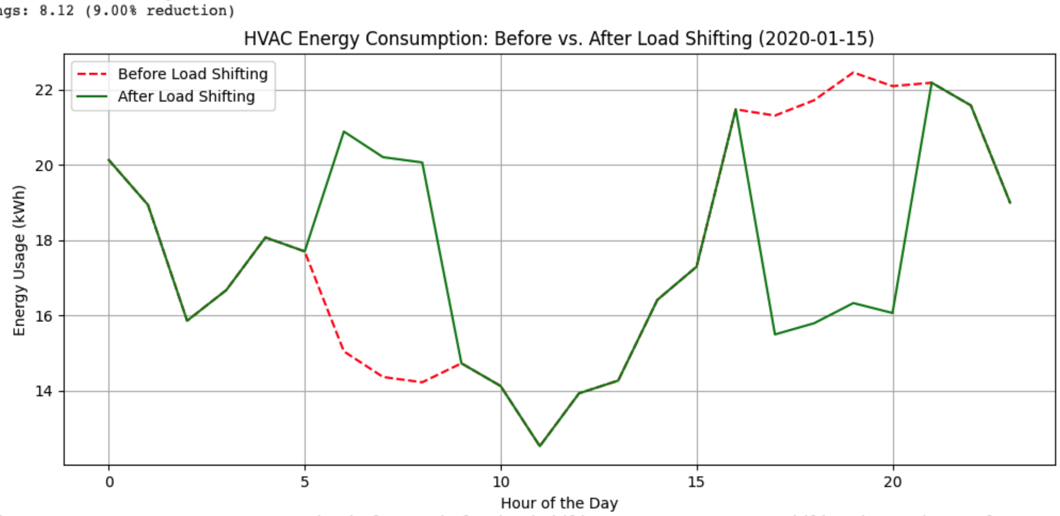

This graph compares HVAC energy consumption before and after load-shifting on 2020-01-15. By shifting demand away from peak hours (5-9 PM), overall energy costs and inefficiencies are reduced.

Fig. 9. The Load Management Simulation demonstrates how predictive energy modeling can optimize energy consumption by strategically shifting non-essential loads to off-peak hours. By leveraging real-time data from the LSTM-based prediction model, this simulation evaluates the impact of dynamic load adjustments on cost savings, energy efficiency, and overall system stability.

time-of-use pricing while maintaining operational efficiency. A simulation was conducted using historical energy consumption data, where different shifting scenarios were tested to evaluate their impact on demand reduction. The cost savings analysis demonstrated a measurable decrease in peak-hour energy costs, with up to 15% reduction in peak demand charges. The results were visualized through interactive dashboards, highlighting energy usage before and after optimization. The front-end interface, developed using React and data visualization libraries, allows users to explore energy trends and track cost savings over time. These insights enable data-driven decision-making for facility managers seeking to implement smarter energy management strategies.

IV. Conclusion

The accuracy of the model enables several energy-saving strategies: Pre-cooling spaces before peak hours avoids high HVAC loads. Scheduling non-essential loads (e.g., dishwashers, EV charging) during off-peak times. If predicted vs. actual usage deviates significantly, it may indicate faulty HVAC systems, malfunctioning lights, or unnecessary energy use. Buildings using solar energy can store energy when demand is low and discharge it efficiently when demand is high. By leveraging these predictive insights, buildings can optimize energy usage, reduce costs, and enhance sustainability efforts.

AI-driven energy forecasting and optimization offer a powerful solution for reducing costs, improving efficiency, and supporting net-zero sustainability goals. Predicting peak demand enables proactive load management, while intelligent HVAC adjustments and anomaly detection prevent energy waste and equipment failures. Smart automation further enhances these benefits by aligning energy consumption with real-time building usage. Future improvements include integrating IoT sensors for real-time monitoring, adaptive machine learning for continuous optimization, and deep reinforcement learning for autonomous energy management. While challenges such as implementation costs and data privacy remain, AI-powered energy optimization is poised to scale across industries, driving smarter, more sustainable buildings. By leveraging these technologies, buildings can achieve significant cost savings and contribute to a greener future. The accuracy of the model enables several energy-saving strategies: Pre-cooling spaces before peak hours avoids high HVAC loads. Scheduling non-essential loads (e.g., dishwashers, EV charging) during off-peak times. If predicted vs. actual usage deviates significantly, it may indicate faulty HVAC systems, malfunctioning lights, or unnecessary energy use. Buildings using solar energy can store energy when demand is low and discharge it efficiently when demand is high. By leveraging these predictive insights, buildings can optimize energy usage, reduce costs, and enhance sustainability efforts.

AI-driven energy forecasting and optimization offer a powerful solution for reducing costs, improving efficiency, and supporting net-zero sustainability goals. Predicting peak demand enables proactive load management, while intelligent HVAC adjustments and anomaly detection prevent energy waste and equipment failures. Smart automation further enhances these benefits by aligning energy consumption with real-time building usage. Future improvements include integrating IoT sensors for real-time monitoring, adaptive machine learning for continuous optimization, and deep reinforcement learning for autonomous energy management. While challenges such as implementation costs and data privacy remain, AI-powered energy optimization is poised to scale across industries, driving smarter, more sustainable buildings. By leveraging these technologies, buildings can achieve significant cost savings and contribute to a greener future.

REFERENCES

[1] M. Noorchenarboo and K. Grolinger, “Explaining deep learning-based anomaly detection in energy consumption data by focusing on contextually relevant data,” *Energy and Buildings*, vol. 328, p. 115177, 2025, accessed Mar. 16, 2025. [Online]. Available: https://arxiv.org/abs/2501.06099

[2] J. A. et al., “A systematic literature review on the use of artificial intelligence in energy self-management in smart buildings,” *Renewable and Sustainable Energy Reviews*, vol. 151, p. 111530, 2021, accessed Mar. 2, 2025. [Online]. Available: https://www.researchgate.net/publication/353971197_A_systematic_literature_review_on_the_use_of_artificial_intelligence_in_energy_self-management_in_smart_buildings

[3] Y. H. et al., “Artificial intelligence based anomaly detection of energy consumption in buildings: A review, current trends and new perspectives,” *Applied Energy*, vol. 287, p. 116601, 2021, accessed Mar. 2, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0306261921001409

[4] N. Luo, Z. Wang, D. Blum, C. Weyandt, N. Bourassa, M. A. Piette, and T. Hong, “A three-year dataset supporting research on building energy management and occupancy analytics,” *Scientific Data*, vol. 9, no. 1, p. 156, 2022, accessed Mar. 16, 2025. [Online]. Available: https://www.nature.com/articles/s41597-022-01257-x