Intelligent Machine-Learning-Driven Load Management for Campus-Scale Energy Optimization and Resilience

Salim Oyinlola [1] **Peter Olabisi Oluseyi** [1]

Abstract

The increasing demand for reliable and efficient energy distribution in educational institutions necessitates the adoption of intelligent energy management systems. This research develops a machine-learning-based framework for load management on a university campus using the University of Lagos as a case study. A dataset comprising 3,648 hourly timestamps across 55 occupied structures was analysed: clustering techniques (K-Means, Hierarchical, Gaussian Mixture Models, Spectral, Mini-Batch K-Means and DBSCAN) were applied, with Mini-Batch K-Means achieving the best performance (Silhouette 0.461, Davies–Bouldin 0.767, Calinski–Harabasz 42.0) and segmenting buildings into high-, medium-, and low-demand clusters. Short-term load forecasting (STLF) was then performed with Prophet, SARIMA, ARIMA, LSTM, and GRU on both the whole-campus series and on cluster-specific series; ARIMA produced the lowest point-forecast errors in all evaluations (whole-campus MAPE 7.2% and RMSE 118.7; cluster-level MAPE 3.8–5.4%), and consumption-based clustering reduced MAPE by roughly 33–36% and RMSE by 76–78% across all algorithms. Overall, the study demonstrates the feasibility of applying machine learning for institutional load management, offering a scalable framework for university campus environments.

1. Introduction

Nigeria's electricity sector faces profound challenges that undermine reliable power delivery across the nation. The national grid, with a nominal installed capacity of approximately 12,500 MW, operates at a significantly reduced functional capacity, struggling to meet a peak demand exceeding 20,000 MW (Iwayemi, 2008; Oyedepo, 2012). This chronic shortfall is exacerbated by aging transmission and distribution infrastructure, inadequate maintenance, and inefficiencies in regulatory oversight, resulting in frequent grid instability. Consequently, load shedding and rotational blackouts are employed as emergency measures to balance supply and demand, disrupting economic and social activities (Reuters, 2024). These disruptions have significant impacts, constraining industrial productivity, inflating operational costs, and undermining critical services in sectors such as healthcare and education (Adenikinju, 2005; World Bank, 2023). At the University of Lagos, a leading federal institution, this unreliable power supply disrupts academic operations, research activities, and student life, posing a significant barrier to institutional efficiency.

The University of Lagos campus microgrid uses grid electricity with diesel-powered generators intermittently to meet the energy demands of diverse facilities, including academic buildings, research laboratories, administrative offices, and residential hostels. The current energy management strategy employs a rigid, time-based load shedding schedule, allocating power to academic and administrative buildings from 8:00 AM to 6:00 PM and to residential areas from 6:00 PM to 8:00 AM. This approach fails to accommodate real-time variations in energy consumption patterns, leading to inefficiencies. Critical facilities such as medical quarters and research laboratories may experience power shortages during peak academic hours, while low-demand areas such as unoccupied lecture halls consume unnecessary power. The absence of building-level metering infrastructure complicates precise demand monitoring, forcing reliance on aggregated campus-wide data or manual estimations based on appliance inventories (Okonkwo & Ntibi, 2015).

To address these challenges, this study proposes a machine-learning-based clustering framework for energy management on the university campus. The system groups campus buildings based on estimated hourly energy consumption patterns, derived from appliance inventories and typical usage schedules, and integrates these clusters into an intelligent load-shedding framework that dynamically ranks buildings by consumption intensity and operational impor-

[1]Department of Electrical and Electronics Engineering, University of Lagos, Akoka, Nigeria. Correspondence to: Salim Oyinlola <salimoyinlola@ieee.org>.

Proceedings of the 6th Muslims in ML Workshop at the 43rd International Conference on Machine Learning, Seoul, South Korea. 2026. *Copyright 2026 by the author(s).*

tance. By forecasting future load demand at the cluster level using ARIMA, SARIMA, LSTM, GRU, and Prophet models, the system enables proactive load-shedding decisions, ensuring that limited electricity is distributed efficiently and with minimal disruption to academic, residential, and administrative operations.

The specific objectives of this research are to: (i) develop a methodology for estimating hourly energy consumption for individual buildings using appliance inventories and typical usage patterns, validated against aggregated campus data; (ii) implement a clustering algorithm with an optimal cluster count to group buildings based on consumption patterns, evaluated using Silhouette, Davies–Bouldin, and Calinski–Harabasz scores; and (iii) develop a short-term load forecasting (STLF) method to predict future load demand for each cluster.

The significance of this study lies in its innovative approach to overcoming data scarcity in the absence of building-level metering. By estimating consumption through appliance inventories, the methodology provides actionable insights into building-specific energy usage patterns, enabling informed load-shedding decisions. The framework is scalable, offering a model for other Nigerian universities and institutions facing similar energy constraints, and aligns with Sustainable Development Goals 7 (Affordable and Clean Energy) and 9 (Industry, Innovation, and Infrastructure).

2. Related Work

The increasing complexity of microgrid systems, particularly in environments with unreliable power grids, has driven the adoption of machine-learning techniques for energy management, load forecasting, and grid optimization. This section synthesises recent advancements in ML applications for energy management, with a focus on clustering, load shedding, and microgrid optimization, and identifies critical gaps relevant to a clustering-based load-shedding system for the University of Lagos.

2.1. Machine Learning in Energy Management

Recent studies have explored ML techniques for energy forecasting and grid management, though many focus on large-scale grids rather than campus-specific microgrids. Aneesh & Wen Ren (2023) employed a convolutional LSTM hybrid model to forecast grid frequency in university microgrids, achieving a 23% reduction in mean absolute error compared to traditional statistical methods. Their approach leverages temporal dependencies in energy consumption data but lacks an operational framework for translating forecasts into load-shedding decisions. Similarly, Mishra et al. (2024) utilized bidirectional LSTM and Transformer models for multi-timescale load prediction in microgrids, demonstrating superior performance in capturing long-term trends; however, their work does not integrate clustering or priority-based load shedding. Soman et al. (2020) combined wavelet transforms with XGBoost to forecast peak demand in microgrids, achieving an 18% reduction in grid stress through proactive load balancing, but the model does not address prioritization of critical loads during shortages. Sathe & Yang (2023) developed a ConvLSTM model to capture spatial and temporal patterns in grid stability but stops short of operational load-shedding protocols, focusing instead on predictive analytics. Zhou et al. (2021) proposed a semi-supervised learning framework for large-scale transmission grids, reducing computational complexity by 40%; however, its focus on high-voltage grids limits applicability to decentralised campus microgrids where building-level granularity is essential.

2.2. Clustering and Load Prioritization

Clustering techniques have been increasingly applied to group buildings or loads based on consumption patterns, enabling targeted energy management strategies. Elhabyb et al. (2024) used K-Means clustering to segment educational buildings by energy usage profiles, identifying high-consumption facilities for targeted efficiency interventions. Their study, however, does not integrate clustering with real-time load shedding, focusing instead on long-term energy planning. Singh et al. (2024) applied hierarchical clustering to group loads in grid-connected microgrids, improving demand-response efficiency. Their approach lacks a mechanism for dynamic load prioritization during power shortages, a critical need in unreliable grid environments like Nigeria's.

2.3. Load Management and Microgrid Optimization

Load management strategies have been investigated to mitigate power shortages, particularly in regions with unreliable grids. Bourhnane et al. (2020) proposed a scheduling-based approach for energy efficiency in microgrids, using rule-based load management to minimise disruptions; their method lacks ML integration, relying on static priorities that cannot adapt to dynamic consumption patterns. Akindeji & Ewim (2023) analysed hybrid power systems for Nigerian universities, integrating diesel generators and solar photovoltaics to reduce grid dependency, but did not employ ML for load prioritization or shedding. Le et al. (2022) developed a load-shedding framework for industrial microgrids using optimization algorithms; however, it is not tailored to the diverse load profiles of university campuses, where academic, administrative, and residential loads require distinct prioritization. Li et al. (2022) proposed a rule-based load-shedding system for smart grids, achieving a 15% reduction in outage duration, but lacks clustering to differentiate load types. Toro-Mendoza et al. (2023) explored priority-based load shedding for microgrids, focusing on industrial appli-

cations with limited relevance to educational institutions.

2.4. Gaps in Knowledge

Despite these advancements, several critical gaps remain. First, most studies focus on energy forecasting or optimization for large-scale grids, neglecting the unique dynamics of university campus microgrids, such as decentralised generation, diverse load profiles, and the need for real-time load prioritization. Second, while clustering has been used to segment loads, few studies integrate clustering with operational load-shedding frameworks, particularly in resource-constrained settings where building-level metering is absent. Third, existing load-shedding strategies often rely on static or rule-based approaches, lacking the adaptability of ML-based clustering to handle dynamic consumption patterns. Finally, computationally intensive models such as reinforcement learning or deep neural networks are impractical for real-time applications in environments with limited computational resources, highlighting the need for lightweight algorithms.

The present study addresses these gaps by developing a clustering-based load-management framework tailored to the University of Lagos campus microgrid, using K-Means-family algorithms to group buildings based on estimated consumption patterns and integrating these clusters into a priority-based load-shedding pipeline. This approach is particularly suited to data-scarce environments, as it relies on appliance-inventory-based estimations rather than real-time metering.

3. Methodology

This study develops a data-driven framework for optimising energy management on a university campus by addressing the challenges of an unreliable power grid and the absence of building-level metering. The methodology employs machine learning, specifically K-Means clustering, to group campus buildings based on estimated hourly energy consumption patterns, enabling a priority-based load-shedding strategy. The methodology is implemented using Python, leveraging libraries such as `pandas`, `numpy`, and `scikit-learn`.

3.1. Data Collection and Preprocessing

The dataset comprises hourly energy consumption estimates for 55 University of Lagos buildings over 3,648 time points (approximately five months at hourly intervals, from January to May 2025). Due to the absence of individual building-level meters, consumption was estimated using appliance inventories (e.g., lighting, HVAC, laboratory equipment) and typical usage schedules, validated against aggregated campus-wide consumption data from the University works department. The dataset is structured as a $3{,}648 \times 55$ matrix, with rows representing hourly timestamps and columns representing buildings, each entry denoting estimated kWh consumption.

To ensure data quality, preprocessing addressed missing values and scale differences:

Missing value imputation. Missing entries, arising from incomplete appliance inventories or recording errors (approximately 2% of the dataset), were filled with the mean consumption of each building's time series to preserve statistical properties.

Normalisation. Z-score normalisation was applied to each building's consumption series to standardise scales across buildings with varying appliance loads:

$$x_{\text{norm}} = \frac{x - \mu}{\sigma} \tag{1}$$

where x is the original consumption value, and μ and σ are the mean and standard deviation of the building's time series. This ensures comparability for clustering by mitigating bias from buildings with higher baseline consumption.

3.2. Clustering

To group buildings by consumption patterns, clustering was performed using K-Means, Mini-Batch K-Means, Agglomerative Clustering, Gaussian Mixture Models, Spectral Clustering, and DBSCAN. The data was transposed to treat buildings as samples (55 samples) and timestamps as features (3,648 features), capturing temporal consumption patterns. K-Means minimises intra-cluster variance with the objective

$$\mathcal{J} = \sum_{i=1}^{k} \sum_{x \in C_i} \|x - \mu_i\|^2, \tag{2}$$

where C_i is the i-th cluster, μ_i is its centroid, and k is the number of clusters.

The optimal cluster count ($k = 3$) was determined using Silhouette, Davies–Bouldin, and Calinski–Harabasz scores. The Silhouette score measures intra-cluster cohesion and inter-cluster separation:

$$s(i) = \frac{b(i) - a(i)}{\max\big(a(i), b(i)\big)}, \tag{3}$$

where $a(i)$ is the average distance of a point to others in its cluster and $b(i)$ is the average distance to points in the nearest neighbouring cluster. The Davies–Bouldin Index measures average similarity between each cluster and its most similar one:

$$\text{DBI} = \frac{1}{k} \sum_{i=1}^{k} \max_{j \neq i} \frac{S_i + S_j}{M_{ij}}, \tag{4}$$

where S_i is the average intra-cluster distance for cluster i and M_{ij} is the distance between centroids of clusters i and j; lower values indicate better separation. The Calinski–Harabasz Index evaluates how well clusters are separated while remaining internally cohesive:

$$\mathrm{CHI} = \frac{\mathrm{Tr}(B_k)}{\mathrm{Tr}(W_k)} \cdot \frac{n-k}{k-1}, \tag{5}$$

where $\mathrm{Tr}(B_k)$ is the trace of the between-cluster scatter matrix and $\mathrm{Tr}(W_k)$ is the trace of the within-cluster scatter matrix.

3.3. Short-Term Load Forecasting (STLF)

Short-term load forecasting constitutes the core analytical component of this research, bridging unsupervised building clustering and intelligent load-shedding optimization. The objective is to predict hourly electricity demand for each consumption cluster and for the entire campus microgrid over a 24–72-hour horizon, enabling proactive scheduling.

The study implements a multi-model forecasting architecture consisting of statistical, probabilistic, and deep-learning paradigms, capturing different structural properties of the load series (seasonality, autocorrelation, nonlinear behaviour, and long-range temporal dependencies). The redundancy supports cross-model validation in the presence of noisy or incomplete data.

ARIMA. The Autoregressive Integrated Moving Average model combines auto-regression on past values, differencing to enforce stationarity, and a moving-average component over past errors:

$$\Big(1-\sum_{i=1}^{p}\phi_i B^i\Big)(1-B)^d y_t = \mu + \Big(1+\sum_{i=1}^{q}\theta_i B^i\Big)\varepsilon_t, \tag{6}$$

where y_t is the observation at time t, B is the backshift operator ($By_t = y_{t-1}$), d is the order of differencing, ϕ_i and θ_i are the AR and MA coefficients, and $\varepsilon_t \sim \mathcal{N}(0, \sigma^2)$.

SARIMA. Seasonal ARIMA extends ARIMA by explicitly modelling seasonal effects, which are highly relevant for campus energy consumption (daily and weekly cycles driven by class schedules and residential activity). SARIMA introduces seasonal terms (P, D, Q, s), where s is the seasonal cycle length (e.g., 24 for daily, 168 for weekly):

$$\Phi(B^s)\phi(B)(1-B)^d(1-B^s)^D y_t = \mu + \Theta(B^s)\theta(B)\varepsilon_t, \tag{7}$$

where Φ and Θ are the seasonal AR and MA polynomials.

LSTM. Long Short-Term Memory networks capture long-term temporal dependencies in sequential data using memory cells and gating mechanisms (forget, input, and output gates). At each timestep t, the gates are

$$f_t = \sigma(W_f x_t + U_f h_{t-1} + b_f), \tag{8}$$
$$i_t = \sigma(W_i x_t + U_i h_{t-1} + b_i), \tag{9}$$
$$o_t = \sigma(W_o x_t + U_o h_{t-1} + b_o), \tag{10}$$

with cell and hidden updates

$$\tilde{c}_t = \tanh(W_c x_t + U_c h_{t-1} + b_c), \tag{11}$$
$$c_t = f_t \odot c_{t-1} + i_t \odot \tilde{c}_t, \tag{12}$$
$$h_t = o_t \odot \tanh(c_t). \tag{13}$$

GRU. The Gated Recurrent Unit is a variant of the RNN similar to LSTM but with a simplified gating mechanism that merges the input and forget gates into a single update gate:

$$z_t = \sigma(W_z x_t + U_z h_{t-1} + b_z), \tag{14}$$
$$r_t = \sigma(W_r x_t + U_r h_{t-1} + b_r), \tag{15}$$
$$\tilde{h}_t = \tanh\big(W_h x_t + U_h(r_t \odot h_{t-1}) + b_h\big), \tag{16}$$
$$h_t = (1-z_t) \odot h_{t-1} + z_t \odot \tilde{h}_t. \tag{17}$$

GRU was tested as a computationally efficient alternative to LSTM.

Prophet. Prophet is an additive model that decomposes a time series into trend, seasonality, and holiday effects:

$$y(t) = g(t) + s(t) + h(t) + \varepsilon_t, \tag{18}$$

where $g(t)$ is the trend (linear or logistic with changepoints), $s(t)$ is the seasonal component (Fourier series), $h(t)$ captures holiday or event effects, and ε_t is the random error. Prophet is robust to missing values, outliers, and irregularities, making it suitable for real-world energy datasets where readings may be incomplete or noisy.

4. Results

This section presents the results obtained from implementing the framework on the University of Lagos campus, which leverages a clustering-based approach to group buildings by estimated hourly energy consumption and subsequently builds predictive models to forecast future total campus energy use. The results are drawn from a dataset of 3,648 hourly timestamps across 55 buildings.

4.1. Clustering Outcomes

The clustering analysis segmented the 55 buildings using consumption profiles over 3,648 hourly time points, represented as a $55 \times 3{,}648$ matrix with buildings as samples and timestamps as features. Six algorithms (K-Means, Hierarchical, Gaussian Mixture Models, Spectral, Mini-Batch K-Means and DBSCAN) were evaluated. The optimal cluster

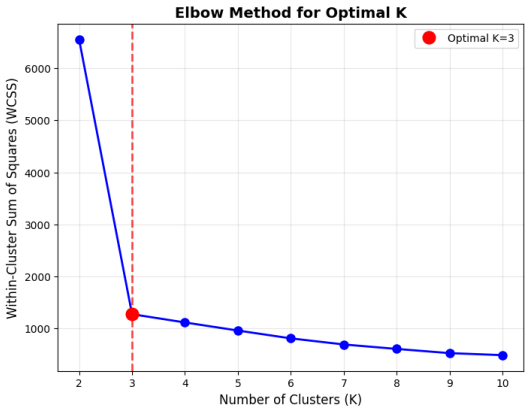

Figure 1. Elbow plot showing within-cluster sum of squares (WCSS) against the number of clusters k, with the optimal $k = 3$ highlighted.

Table 1. Performance metrics of clustering algorithms at $k = 3$.

Algorithm	Silhouette ↑	DB ↓	CH ↑
K-Means	0.4276	0.8315	42.1
Hierarchical	0.4358	0.8251	42.6
GMM	0.4125	0.8453	41.8
Spectral	0.4761	0.7670	40.6
Mini-Batch K-Means	0.4607	0.8000	42.0
DBSCAN	0.1067	2.0661	4.7

count of three was determined through the Elbow Method, in which the within-cluster sum of squares against the number of clusters k identifies the point beyond which adding more clusters yields negligible improvement (Figure 1). Performance metrics for all algorithms are reported in Table 1.

Among the tested clustering algorithms, Mini-Batch K-Means achieved the best overall performance, with a Silhouette score of 0.461, a Davies–Bouldin Index of 0.800, and a Calinski–Harabasz Index of 42.0. Based on these metrics, the dataset was partitioned into three distinct clusters:

- **Cluster 0** (9 buildings): primarily high-consumption academic and administrative buildings such as Engineering, Senate House, and the Auditorium Chiller.
- **Cluster 1** (37 buildings): the largest group, containing student hostels, libraries, and residential quarters.
- **Cluster 2** (9 buildings): specialised facilities such as the Health Centre, Shopping Complexes, and places of worship.

Figure 2 compares the algorithms across the three metrics. The Davies–Bouldin Index, which is lower when clusters are well-separated and internally compact, was minimised by Spectral clustering, while Mini-Batch K-Means achieved the highest Calinski–Harabasz Index. No single algorithm dominated across all three criteria, but Mini-Batch K-Means

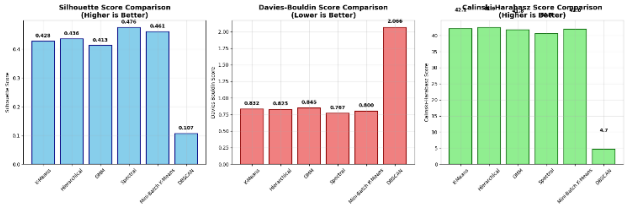

Figure 2. Comparative performance of clustering algorithms across the three evaluation metrics: Silhouette (higher is better), Davies–Bouldin (lower is better), and Calinski–Harabasz (higher is better).

Table 2. Forecasting performance: Mean Absolute Percentage Error (MAPE, %) on the whole campus and per cluster.

Algorithm	Whole	Cluster 0	Cluster 1	Cluster 2
ARIMA	**7.2**	**3.8**	**4.6**	**5.4**
Prophet	8.0	4.1	5.3	6.1
SARIMA	8.4	4.4	5.7	6.6
LSTM	9.1	4.7	6.2	7.5
GRU	11.5	5.2	7.8	9.2

was the best on average and was selected as the clustering algorithm for the downstream forecasting pipeline. DBSCAN underperformed due to its sensitivity to parameter selection at this dataset scale.

4.2. Short-Term Load Forecasting Results

Following clustering, short-term load forecasting was performed at the cluster level. The objective was to predict hourly load profiles 24 hours ahead. The dataset for each cluster was partitioned into training (80%) and testing (20%) subsets, yielding approximately 2,918 training and 730 testing samples per cluster. Tables 2 and 3 report MAPE and RMSE for the five forecasting algorithms on the whole-campus series and on each cluster. Figure 3 visualises the RMSE comparison across models and forecasting approaches, and Figure 4 compares per-cluster MAE, RMSE, and MAPE across algorithms.

ARIMA attains the lowest error in every reported cell of Tables 2 and 3: whole-campus MAPE $= 7.2\%$ and RMSE $= 118.7$; Cluster 0 MAPE $= 3.8\%$ and RMSE $= 21.8$; Cluster 1 MAPE $= 4.6\%$ and RMSE $= 44.7$; Cluster 2 MAPE $= 5.4\%$ and RMSE $= 13.0$.

Considering whole-campus results, ARIMA reduces MAPE by 0.8 pp (10.0% relative) and RMSE by 10.2 units (7.9% relative) compared with Prophet; by 1.2 pp (14.3%) and 15.5 units (11.6%) compared with SARIMA; by 1.9 pp (20.9%) and 24.1 units (16.9%) compared with LSTM; and by 4.3 pp (37.4%) and 48.7 units (29.1%) compared with GRU.

The effect of modelling at the cluster level is quantitatively large for every algorithm. For ARIMA, the mean of

Table 3. Forecasting performance: Root Mean Square Error (RMSE) on the whole campus and per cluster.

Algorithm	Whole	Cluster 0	Cluster 1	Cluster 2
ARIMA	**118.7**	**21.8**	**44.7**	**13.0**
Prophet	128.9	23.2	49.8	15.6
SARIMA	134.2	24.1	52.9	16.8
LSTM	142.8	24.6	58.3	19.7
GRU	167.4	26.4	67.8	23.9

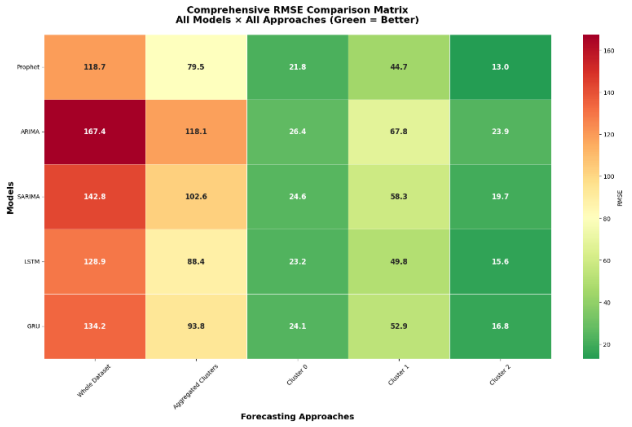

Figure 3. RMSE comparison matrix across all five models and four forecasting approaches (whole dataset, aggregated clusters, and Clusters 0–2). Greener cells indicate lower error.

the three cluster MAPEs is 4.6% versus the whole-dataset MAPE of 7.2% (an absolute reduction of 2.6 pp, 36.1% relative); the mean cluster RMSE is 26.5 versus the whole RMSE of 118.7 (-92.2 units, 77.7% relative). Equivalent cluster-vs-whole reductions for the other methods are: GRU mean cluster MAPE 7.4% vs. whole 11.5% (-35.7%) and RMSE 39.4 vs. 167.4 (-76.5%); LSTM mean cluster MAPE 6.13% vs. 9.1% (-32.6%) and RMSE 34.2 vs. 142.8 (-76.1%); Prophet mean cluster MAPE 5.17% vs. 8.0% (-35.4%) and RMSE 29.53 vs. 128.9 (-77.1%); SARIMA mean cluster MAPE 5.57% vs. 8.4% (-33.7%) and RMSE 31.27 vs. 134.2 (-76.7%). Across algorithms the reduction in mean cluster MAPE ranges from 32.6% to 36.1%, and the reduction in mean cluster RMSE ranges from 76.1% to 77.7%.

Aggregating these observations, ARIMA is the lowest-error model in all eight reported evaluations (four MAPE cells and four RMSE cells). The two dominant findings are (1) ARIMA produces the lowest MAPE and RMSE in every reported case, and (2) segmentation into three consumption-based clusters reduces MAPE by roughly 33–36% and RMSE by roughly 76–78% across the candidate algorithms. Figure 5 visualises the predicted vs. actual relationship for the whole-campus ARIMA forecast, showing a trend line close to the identity ($y \approx 1.05x$) consistent with low systematic bias.

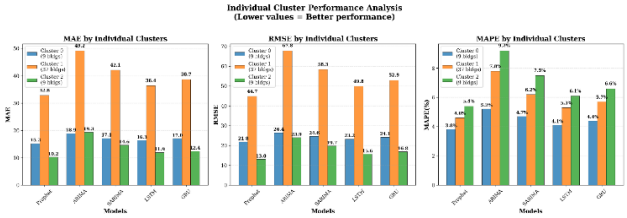

Figure 4. Individual cluster performance analysis: MAE, RMSE, and MAPE by algorithm for each of the three clusters.

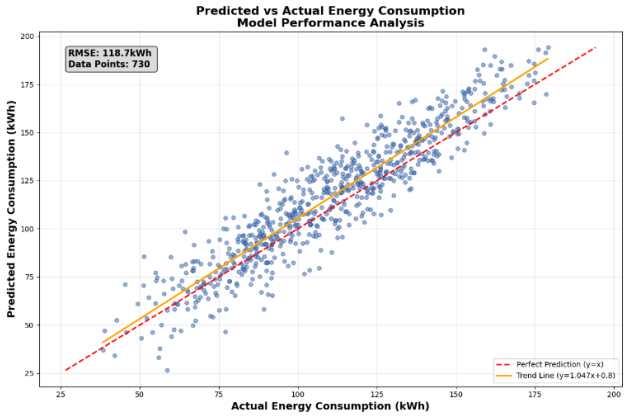

Figure 5. Predicted vs. actual hourly energy consumption (kWh) for the whole-campus ARIMA forecast on the held-out test set (RMSE $= 118.7$, $n = 730$). The dashed line shows the ideal $y = x$ relationship; the trend line is $y = 1.047x + 0.8$.

These quantitative results should be used to guide model selection and operational deployment: ARIMA provides the most accurate point forecasts in this evaluation, and applying cluster-specific models yields large, quantifiable improvements in both relative (MAPE) and absolute (RMSE) error.

5. Conclusion

This study developed and evaluated a machine-learning-based framework intended to improve load management at the University of Lagos by providing building-level energy visibility and data-driven control strategies in an environment lacking pervasive submetering. A reproducible methodology was derived to estimate hourly energy consumption at the appliance and building levels using building-specific appliance inventories and typical usage patterns; the formulation enables generation of granular hourly load profiles where direct metering is absent. Using those estimated profiles, an unsupervised clustering pipeline grouped buildings according to consumption signatures rather than geography, producing consumption-centric categories such as high-load laboratories, moderate-use academic blocks, and low-use administrative or residential blocks. Predictive models for hourly demand (ARIMA, SARIMA, LSTM,

GRU, and Prophet) were then evaluated on each cluster, with ARIMA producing the lowest MAPE and RMSE across both the whole-campus and the per-cluster evaluations, and consumption-based clustering yielding 33–36% MAPE reductions and 76–78% RMSE reductions across all candidate algorithms.

The framework addresses a critical operational gap, the absence of per-building energy visibility, and shows potential to reduce operational inefficiencies, support targeted conservation, and enable more equitable demand-side interventions. It is consistent with broader sustainability goals (SDGs 7, 11, and 13) by promoting responsible energy consumption and facilitating adaptive, intelligent energy systems. Key limitations include the reliance on survey-derived usage patterns, which introduce estimation error relative to direct metering; the need for higher-resolution validation data to refine and calibrate models; and the absence of a full live deployment to validate real-world operational performance. To translate the framework from research into operational benefit, future work should pilot the approach on a representative subset of buildings (at least one high-load laboratory, one academic block, and one residential or administrative facility), prioritise submetering and environmental sensors on those buildings to obtain ground-truth measurements, and integrate the consumption estimator, clustering module, and forecasting engine into a modular operator-facing decision-support system with fail-safe mechanisms for critical loads.

Impact Statement

This work uses appliance-inventory-based estimates of building-level energy consumption from a Nigerian university campus, with no customer-level identifiers, demographic attributes, or other personally identifying information. The proposed framework is intended to support more efficient, equitable load management in resource-constrained institutional microgrids; potential beneficial impacts include reduced diesel-generator dependence and more predictable power availability for critical facilities (laboratories, medical quarters). We acknowledge that prioritised load shedding has fairness implications across building populations (academic vs. residential, central vs. peripheral); any operational deployment should therefore explicitly characterise per-building service-level outcomes and be subject to oversight by an institutional energy management committee with stakeholder representation. The estimation methodology rests on survey-derived usage patterns and should be re-calibrated against submetered ground truth before any consequential operational decision is automated.

References

Adenikinju, A. Analysis of the cost of infrastructure failures in a developing economy: The case of electricity sector in Nigeria. Technical report, African Economic Research Consortium, 2005.

Akindeji, K. T. and Ewim, D. R. E. Economic and environmental analysis of a grid-connected hybrid power system for a University Campus. *Bulletin of the National Research Centre*, 47(1):75, 2023.

Aneesh, S. and Wen Ren, Y. Grid frequency forecasting in university campuses using convolutional LSTM. *arXiv preprint arXiv:2310.16071*, 2023.

Bourhnane, S., Abid, M. R., Lghoul, R., Zine-Dine, K., and Benhaddou, D. Machine learning for energy consumption prediction and scheduling in smart buildings. *SN Applied Sciences*, 2(2):297, 2020.

Elhabyb, K., Baina, A., Bellafkih, M., and Deifalla, A. F. Machine learning algorithms for predicting energy consumption in educational buildings. *International Journal of Energy Research*, 2024:6812425, 2024.

Iwayemi, A. Investment in electricity generation and transmission in Nigeria: Issues and options. International Association for Energy Economics, 2008. URL https://www.iaee.org/documents/newsletterarticles/Iwayemi.pdf.

Le, N. T., Phung, T. T., Quyen, A. H., Nguyen, B. L. P., and Nguyen, A. N. A hybrid approach of artificial neural network-particle swarm optimization algorithm for optimal load shedding strategy. *International Journal of Electrical and Computer Engineering (IJECE)*, 12(4): 4253–4263, 2022.

Li, J., Chen, S., Wang, X., and Pu, T. Load shedding control strategy in power grid emergency state based on deep reinforcement learning. *CSEE Journal of Power and Energy Systems*, 8(4):1175–1182, 2022.

Mishra, S., Frank, S. M., Petersen, A., Buechler, R., and Slovensky, M. Deep-learning-based multi-timescale load forecasting in buildings: opportunities and challenges from research to deployment. *Future Energy*, 3(4):47–55, 2024.

Okonkwo, C. and Ntibi, J. Exploiting renewable energy for rural development in Nigeria. *Niger Delta Journal of Engineering and Technology*, 2(1):190–197, 2015.

Oyedepo, S. O. Energy and sustainable development in Nigeria: the way forward. *Energy, Sustainability and Society*, 2(15), 2012.

Reuters. Why Nigeria's power grid is failing, December 2024. URL https://www.reuters.com/world/africa/why-nigerias-power-grid-is-failing-2024-12-11/.

Sathe, A. and Yang, R. Grid frequency forecasting in university campuses using convolutional LSTM. *arXiv preprint arXiv:2310.16071*, 2023.

Singh, A. R., Kumar, R. S., Bajaj, M., Khadse, C. B., and Zaitsev, I. Machine learning-based energy management and power forecasting in grid-connected microgrids with multiple distributed energy sources. *Scientific Reports*, 14(1):19207, 2024.

Soman, A., Trivedi, A., Irwin, D., Kosanovic, B., McDaniel, B., and Shenoy, P. Peak forecasting for battery-based energy optimizations in campus microgrids. *arXiv preprint arXiv:2005.13517*, 2020.

Toro-Mendoza, M. A., Segundo-Ramírez, J., Esparza-Gurrola, A., Visairo-Cruz, N., Núñez Gutiérrez, C. A., and Pérez-Negrón, C. Toward adaptive load shedding remedial action schemes in modern electrical power systems. *IEEE Access*, 11:109324–109341, 2023.

World Bank. Nigeria power sector reforms: Fact sheet. Technical report, World Bank, 2023.

Zhou, Y., Park, J., and Zhu, H. Scalable learning for optimal load shedding under power grid emergency operations. *arXiv preprint arXiv:2111.11980*, 2021.