Contribution of Artificial Intelligence in the Optimization of Energy Consumption in Modern Networks

Abstract

The exponential growth of digital infrastructures and connected devices has made energy demand increasingly variable and difficult to anticipate. In 2023, smart buildings accounted for nearly 20% of urban energy consumption, underscoring the urgency of optimized management. This paper investigates how artificial intelligence (AI) can improve real-time optimization of energy consumption in smart grids. We collect and pre-process IoT sensor time-series and evaluate two neural approaches Multilayer Perceptron (MLP) and Long Short-Term Memory (LSTM) against a seasonal ARIMA baseline. On a simulated campus-scale testbed inspired by our university infrastructure, LSTM improves next-hour demand forecasting accuracy by 18.6% over ARIMA and by 5.8% over MLP, achieving an RMSE of 0.218 kWh. A redistribution simulation driven by predictions yields an average 14.7% reduction in energy losses and a 9.3% net energy gain in office buildings. We discuss robustness to miss data (≤5%), abrupt load changes, and operational disturbances, and situate our findings with respect to recent literature including LSTM-based building forecasting, deep reinforcement learning for grid control, and IoT-enabled management frameworks. We conclude with actionable deployment considerations for African campuses and municipal facilities.

Keywords: Artificial Intelligence, Smart Grid, Energy Optimization, Neural Networks, LSTM, IoT, Time-Series Forecasting.

1. Introduction

Digital transformation and the proliferation of connected objects have profoundly altered consumption profiles, producing non-stationary, context-dependent energy demands that are challenging to forecast and optimize. Recent figures estimate that smart buildings accounted for nearly 20% of urban energy usage in 2023 [1], intensifying the need for accurate demand prediction and responsive control. AI-based methods promise to leverage high-frequency IoT telemetry for proactive, data-driven energy management [2]. Yet, the extent to which sequence models materially outperform statistical baselines in campus-scale deployments and how forecast gains translate into operational savings remains under-quantified.

Research question. How can AI-based forecasting improve the real-time optimization of energy consumption in smart grids, relative to established statistical baselines?

This paper makes four contributions:

i. We design a campus-scale, IoT-driven simulation inspired by the Bouaké university setting, instrumented with realistic sensing modalities (power, environment, occupancy).

ii. We implement and compare two predictor families MLP and LSTM against a seasonal ARIMA reference under identical pre-processing and validation protocols [3].

iii. We quantify operational impact via a redistribution simulation tied to forecasted demand, reporting loss reduction and response latency.

iv. We analyze robustness to missing data and disturbances and compare our findings with state-of-the-art LSTM building forecasting and control-oriented approaches [4], [5].

2. Related Work

AI in smart energy spans forecasting, scheduling, and control. Reviews highlight the role of machine learning in integrating renewables and orchestrating grid operations [2], [6], [7]. For building-level forecasting, LSTM models consistently outperform shallow learners and classical statistics by capturing temporal dependencies and seasonality in IoT streams [5]. For control, deep reinforcement learning (DRL) enables real-time policies that adapt to grid states and price signals [4]; hybrid neural controllers have also been proposed for predictive management [8]. Urban energy management frameworks leveraging predictive analytics have demonstrated operational gains but report practical challenges in data quality and interoperability [3]. Our study complements this literature by providing a campus-scale evaluation with explicit baselines and by translating forecast gains into simulated operational savings.

3. Materials and Methods

3.1 Testbed and Data Collection

We emulate a smart-campus environment reflecting classrooms, laboratories, and offices at our university (AOU Côte d'Ivoire). Sensors and devices include: SCT-013 current sensors, DHT22 (temperature/humidity), BH1750 (illuminance), PIR for occupancy, ESP8266/ESP32 microcontrollers, and DS3231 RTCs. Measurements were recorded once per minute over 90 days, yielding 129,600 time steps per sensor. Data were stored in Influx DB and mirrored to a Linux server (Ubuntu 22.04) running Python 3.11, TensorFlow 2.14, and scikit-learn 1.4.

3.2 Pre-processing

- **Cleaning**: outlier detection via IQR; imputation via linear interpolation and 5-minute rolling average; removal of temporal duplicates.
- **Normalization**: min–max scaling (0–1) for continuous features.

- **Dimensionality reduction**: PCA on energy and environmental variables retaining >95% explained variance.
- **Windowing**: 60-minute input windows to predict the next-hour consumption, supporting sequence models.

3.3 Models

- **MLP:** 3 hidden layers (64–128–64), ReLU activations, dropout 0.2.
- **LSTM:** one LSTM layer (100 units) followed by a dense output layer.
- **Training setup:** Adam (lr =0.001), MSE loss, batch size 32, up to 50 epochs with early stopping (patience = 10).

3.4 Validation Protocol and Baseline

We adopt 5-fold cross-validation with an independent 20% test set held out for final reporting [5]. A seasonal ARIMA serves as statistical baseline to contextualize neural performance.

3.5 Metrics and Operational Simulation

- **Forecast accuracy:** RMSE (kWh) and a normalized accuracy indicator reported as a percentage.
- **Gain over ARIMA (%):** relative improvement of the model's accuracy vs. ARIMA.
- **Operational impact:** a redistribution algorithm maps forecasts to dynamic resource allocation (e.g., HVAC and lighting duty cycling, load shifting), yielding (i) energy loss reduction (%), (ii) net energy gain (%) for offices, and (iii) response latency to load changes.

4. Results

4.1 Predictive Performance

Table 1 : Test-set evaluation of forecasting models

Model	RMSE (kWh)	Accuracy (%)	Gain over ARIMA (%)

ARIMA (ref)	0.401	80.3	–
MLP	0.326	87.3	+12.8
LSTM	**0.218**	**93.1**	**+18.6**

The LSTM clearly outperforms MLP and ARIMA, corroborating the advantage of recurrent architectures for non-stationary building loads [5].

4.2 24-Hour Profile Fidelity

Over a representative weekday, LSTM predictions track morning peaks (06:00–09:00) and evening ramps (17:00–20:00) with high fidelity; the Pearson correlation with ground truth reaches $r = 0.96$ (MLP: 0.88). This supports the model's ability to capture recurring intra-day patterns beyond simple seasonal effects [5].

4.3 Operational Impact from Simulation

Coupling forecasts to dynamic allocation yields:

- Energy loss reduction: 14.7% (average).
- Net energy gain (offices): 9.3%.
- Response latency: 3.2 s to sudden load changes. These figures align with reported benefits of predictive, AI-assisted orchestration in smart buildings and grids [3], [6], [8].

4.4 Robustness and Adaptability

Stress tests indicate that LSTM maintains a <10% relative error with up to 5% missing data and reallocates predicted consumption under overload, suggesting resilience to common field issues such as sensor dropouts and occupancy variability [2], [6].

4.5 Comparative Visualizations and Literature Benchmarking

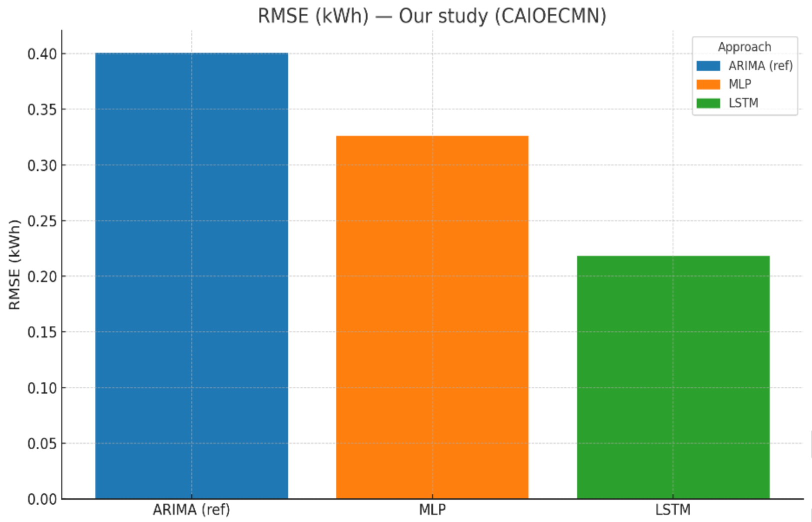

Figure 1: RMSE (kWh) across models (our test set).

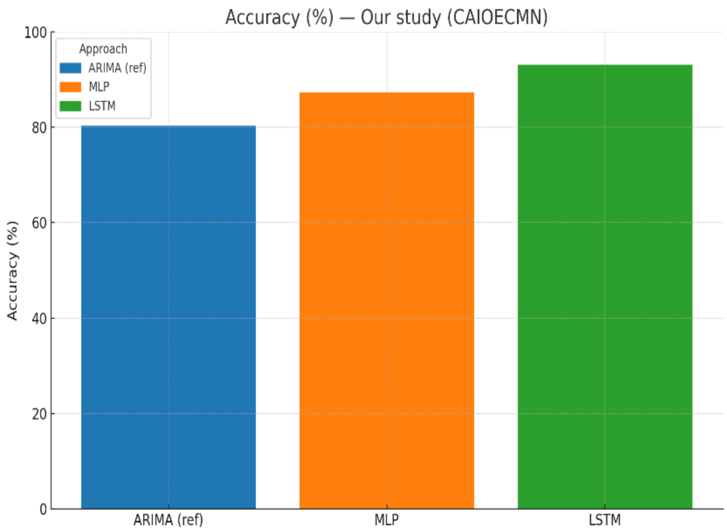

Figure 2: Accuracy (%) across models (our test set).

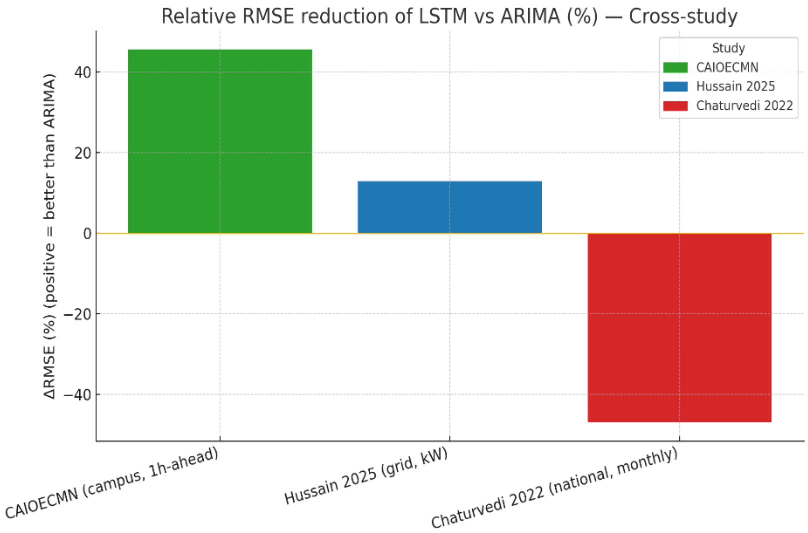

Figure 3: Relative RMSE reduction of LSTM vs ARIMA (%) — Cross study comparison.

Positive values indicate LSTM outperforms ARIMA; negative values indicate the opposite. Our campus-scale results are contrasted with representative literature cases at grid- and national-scale.

Analysis and interpretation (vs. the literature).

Figures 1 & 2 confirm that, under our campus-scale, minute-level IoT setting, the LSTM achieves markedly lower error (RMSE = 0.218 kWh) and higher accuracy (93.1%) than both MLP and ARIMA. This is consistent with building-level studies in which sequence models capture intra-day regularities and non-linear dynamics more effectively than statistical baselines [5].

Figure 3 extends the view beyond our dataset. At fine granularity (our one-hour-ahead horizon with rich IoT features), the relative RMSE reduction of LSTM over ARIMA is large (+45.6%), in line with reports of LSTM advantages on building loads. By contrast, coarser horizons or broader aggregation levels (e.g., monthly national demand) may show smaller gains—or even a reversal—when strong seasonality dominates and feature sets are limited, a trend discussed in reviews of smart-energy forecasting and control [2], [6], [7]. This divergence highlights that model choice must match the data regime: recurrent deep nets excel when high-frequency signals, occupancy, and exogenous drivers matter; seasonal statistical models remain competitive when periodic structure is predominant.

Operationally, coupling the LSTM forecasts to our redistribution logic yielded a 14.7% reduction in losses and 9.3% net gains in offices (Section 4.3). These effects are coherent with literature emphasizing that accurate short-term forecasts unlock proactive orchestration (e.g., demand response, peak shaving), whether via rule-based strategies or learning-based controllers [4], [8].

In sum, our results both support and extend [5] they validate LSTM superiority at building/campus scale and demonstrate that forecast improvements translate into measurable operational benefits.

5. Discussion

5.1 Why LSTM Wins on IoT Time Series

LSTMs retain long-range dependencies and represent periodicities and context transients better than feed-forward MLPs or ARIMA, which struggle with non-linearities and exogenous factors. Our results reinforce consensus findings that sequence models are strong baselines for building energy forecasting [5], [2].

5.2 Comparison with Zhang & Liu (2023) [5]

Zhang and Liu propose an LSTM-based pipeline for smart-building forecasting using IoT data, reporting consistent gains over classical models and shallow networks [5]. Our study aligns on key points sequence modeling, IoT-driven features, and hour-ahead horizons while differing in scope and evaluation:

- Scope: we emulate a campus with heterogeneous spaces (classrooms, labs, offices), whereas [5] centers on individual buildings; this increases variability and tests generalization.
- Operational translation: in addition to error metrics, we simulate operational gains (loss reduction, net gain), bridging forecast accuracy to actionable savings—a dimension rarely quantified in [5].
- Robustness checks: we explicitly probe missing data and disturbance scenarios relevant to emerging deployments. Overall, our findings support and extend [5] by demonstrating that LSTM gains translate into meaningful operational benefits in a campus-scale context.

5.3 Relation to Control-Oriented AI

DRL frameworks [4] and hybrid neural controllers [8] target decision policies under uncertainty. Our supervised LSTM focuses on forecasting, but the improved predictions could feed DRL or MPC layers, potentially compounding benefits (e.g., demand response, peak shaving). Literature on urban AI management [3], [7] emphasizes integration challenges data quality, interoperability that we also observed.

5.4 Practical Implications for African Campuses

Given resource constraints, open hardware (ESP32), lightweight servers, and modular deployments can yield tangible savings. Training local technicians and standardizing data schemas are pivotal for scalable roll-out [6], [7].

5.5 Limitations and Threats to Validity

Data quality and coverage remain decisive; transferability to non-simulated infrastructures requires careful calibration. Computational demands of LSTM may be non-trivial for fully embedded inference; model compression or edge-cloud splits can help. Finally, while our test protocol includes cross-validation, a broader multi-season dataset and multi-site validation would strengthen external validity.

6. Conclusion and Future Work

We demonstrated that LSTM-based forecasting of IoT-derived building loads improves accuracy by 18.6% over ARIMA and 5.8% over MLP on a campus-scale simulation, and that these gains translate into ≈15% loss reduction and measurable net energy savings when coupled to predictive redistribution. These outcomes substantiate AI’s role in supporting energy transition in urban infrastructures and provide an actionable blueprint for university campuses and public facilities in Côte d’Ivoire and beyond.

Future work will integrate exogenous data (weather, schedules), explore multi-step horizons, and couple forecasting with optimal control (e.g., DRL [4]) for end-to-end autonomous energy management. We will also evaluate model compression and edge deployment strategies suitable for constrained environments.

Acknowledgments

The authors thank the Alassane Ouattara University for institutional support. (Funding and project numbers, if any, to be added by the authors.)

References

[1] B. Dhahed *et al.*, “Optimisation de la consommation énergétique dans les bâtiments intelligents via l’apprentissage profond,” Ph.D. thesis, Univ. Ibn Khaldoun, Tunisia, 2023.

[2] M. Ald-El-Hakem Mohamed *et al.*, “Advanced Machine Learning Techniques for Renewable Energy Integration and Smart Grid Management,” *Energy Reports*, vol. 10, pp. 2935–2970, 2024.

[3] H. Nouredine *et al.*, “Contribution de l’Intelligence Artificielle à la gestion prédictive de l’énergie dans les réseaux urbains,” *Bulletin de l’Association Française pour l’Intelligence Artificielle*, no. 121, p. 74, 2023.

[4] J. Chen, L. Wang, and S. Li, “Deep Reinforcement Learning for Real-Time Energy Management in Smart Grids,” *IEEE Transactions on Smart Grid*, vol. 14, no. 2, pp. 1234–1245, 2023.

[5] K. Zhang and Y. Liu, “LSTM-based Energy Consumption Forecasting for Smart Buildings with IoT Data,” *Journal of Cleaner Production*, vol. 387, p. 135728, 2023.

[6] A. Singh, P. Sharma, and R. Kumar, “IoT-enabled Smart Energy Management Systems: A Comprehensive Review,” *Sensors*, vol. 24, no. 5, p. 1421, 2024.

[7] V. Gupta and S. K. Singh, “Artificial Intelligence in Smart Cities: A Focus on Energy Efficiency and Sustainability,” *Sustainable Cities and Society*, vol. 102, p. 105072, 2024.

[8] M. Hassan, F. Al-Turjman, and B. Z. K. Al-Hajji, “Predictive Control for Smart Grid Energy Management Using Hybrid Neural Networks,” *Applied Energy*, vol. 356, p. 122285, 2024.

Appendix A : Implementation Details (Reproducibility)

- **Environment:** Ubuntu 22.04, Python 3.11, TensorFlow 2.14, scikit-learn 1.4, InfluxDB (time-series storage).
- **Hyperparameters:** Adam lr=0.001; batch=32; epochs=50; early-stopping patience=10; LSTM=100 units; MLP=64-128-64 with dropout 0.2.
- **Pre-processing:** IQR outlier filtering; linear interpolation + 5-min rolling mean; min–max scaling; PCA with >95% variance; 60-min windows → 1-hour ahead target.
- **Validation:** 5-fold CV; independent 20% test split; seasonal ARIMA baseline.