© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723105

Intelligent Predictive Energy Management Of 5G Base Stations Using Reinforcement Learning abstract

ONWUHA, U.H
Department of Computer Science, Federal College of Forest Resources Management, Ishiagu, Ebony State, Nigeria.

Abstract-** The aim of the paper is design of AI-based predictive energy management model for efficient wireless base stations using Reinforcement Learning (RL). The methodology employs a multi-phase methodology to develop and validate a reinforcement learning-based predictive energy management model for 5G base stations. Energy consumption data were collected and used to model the base station's operation and formulate an optimization problem. A reinforcement learning algorithm was then implemented in Python to learn real-time energy-saving strategies based on traffic load and environmental factors. The model was validated through simulations and real-time tests, evaluating improvements in energy efficiency, service quality, and cost reduction. The testbed is MTN 5G site at Polo Park, Enugu. Findings from investigation revealed the cost of daily energy used by the site is ₦31377.7787 and the monthly running cost is ₦972711.1397. To reduce the economic burden, RL-based predictive energy management model was proposed as predictive and adaptive learning techniques which utilized Q- learning algorithm as the learning agent to adjust the action of the site based on dynamic state space of the network condition. Python programming language was used to train the model and also integration on the 5G network. Results obtained revealed that with RL, the daily energy saved when compared with characterized is ₦4625.0896. The monthly energy saved amount to ₦170130.4667 with RL-based predictive energy management model, while every year, ₦2041565.6004 is the amount saved. The reason was because of the ability of the RL to adjust the action of RL based on the different state space of the site to save energy. The percentage reduction in daily energy consumed with RL based station is 33.90%. The percentage reduction in monthly energy consumed is 14.74%. In conclusion the study has demonstrated the RL can help manage energy consumption in the 5G network while maintaining quality of service.

***Keywords:** energy management; base stations; reinforcement learning (rl); energy consumption; polo park, enugu*

I. INTRODUCTION

The rapid growth in mobile data traffic and the deployment of 5G and beyond wireless networks have led to significant increases in energy consumption in wireless communication systems, particularly in base stations (BSs) (Wang, 2020). These base stations consume substantial power, contributing to operational costs and environmental impacts through carbon emissions (Ma et al., 2022). Consequently, there is a growing need for energy-efficient solutions in the design and operation of wireless base stations. Traditional energy-saving methods, such as sleep mode techniques and network densification, have provided incremental improvements but are often limited in adapting to dynamic traffic and channel conditions (Ichimescu et al., 2024; Kaur et al., 2023).

Artificial Intelligence (AI) has emerged as a transformative technology with the potential to optimize energy consumption in wireless networks. AI-based models can learn complex patterns in data and make real-time decisions that adapt to changing conditions (Zhu et al., 2025). Specifically, RL, a subset of AI, has demonstrated significant potential in dynamic resource allocation, power control, and energy management. RL enables the system to learn optimal energy-saving strategies through continuous interaction with the network environment, providing a more adaptive and intelligent approach compared to traditional methods (Ichimescu et al., 2024).

In recent years, various studies (Premalatha and SahayaAnselin, 2023; Ezzeddine et al., 2024; Zhu et al., 2025) have applied RL to optimize different aspects of wireless communication networks, such as resource allocation, interference management, and handover processes. These studies have shown promising results in improving network performance while reducing energy consumption. However, many of these works focus on isolated components of the network rather than comprehensive energy management across the entire base station operation (Ma et al., 2023). Moreover, Ezzedddine et al., (2024) revealed that the majority of existing approaches are tailored for static scenarios, limiting their applicability in highly dynamic environments typical of modern wireless networks. By dynamically adjusting transmission power, switching BSs on/off based on traffic loads, and optimizing resource allocation, RL can achieve intelligent energy-saving strategies without compromising Quality of Service (QoS). Such adaptability is crucial in countries like Nigeria, where urban centres face not only rapidly changing user loads but also challenges with grid reliability and environmental concerns tied to carbon-intensive energy generation.

In Nigeria, the energy management challenge is twofold: network operators must ensure reliable and high-quality connectivity to meet the country's digital growth while minimizing energy consumption in a context where the national grid is often unreliable and diesel-powered generators are used extensively for backup. This paper proposes to address these questions through the development of an AI-based predictive energy management model using reinforcement learning tailored for wireless BSs in Nigeria. By applying RL's ability to learn from experience and adapt to environmental changes, this research aims to bridge the gap between traditional static energy management techniques and the dynamic demands of modern wireless networks in emerging economies.

II. METHODOLOGY

The methodology for this study involves a multi-phase approach to develop and validate a reinforcement learning-based predictive energy management model for 5G wireless base stations. First, daily measurements of energy consumed by the base station were conducted, alongside evaluations of cost analysis to establish baseline performance as a means of data collection. Next, a reinforcement learning algorithm was then be designed to dynamically manage and reduce energy usage by learning optimal control strategies in real-time based on traffic load and environmental conditions. Finally, the model will be implemented using Python, with experimental validation carried out through simulations and real-time tests, measuring improvements in energy efficiency, QoS, and operational cost reduction.

Data Collection Through Investigating the Energy Consumption Level of 5G Network Cell

To investigate the energy consumption level of a MTN 5G network cell, a field-based experimental setup was carried out within an operational 5G macrocell environment located in Enugu Urban, Nigeria. The measurement focused on acquiring real-time power usage data over a continuous 30-day cycle to capture daily, weekly, and peak/off-peak variations in energy demand. The primary instrumentation used included a Fluke 1730 Three-Phase Energy Logger, which was installed at the power input terminals of the 5G base station to monitor and log voltage (V), current (A), power factor, active power (kW), and energy consumed (kWh) in real-time. This data was then compressed summing up the total daily power consumed in kW/day and reported in Table 4.1 for analysis. The purpose of this investigation was to establish an empirical baseline of energy demand under varying operational conditions and user loads. This baseline is essential to understand how energy is consumed across different sub-systems of the base station (e.g., baseband units, radio units, cooling units) and the cost implications. The unit cost of electricity used is ₦209.5/kWh. The site considered is polo park base station with cell ID T6320. The site was selected due to the high population density of users within the area. This thorough investigation provided foundational insight for the formulation of the energy optimization problem, enabling the later development of a reinforcement learning-based predictive model that dynamically manages energy based on learned patterns and predicted user demand.

Design of the Reinforcement Learning-Based Predictive Energy Management Model That Dynamically Adjusts Energy Consumption in Wireless Base Stations.

Traditional energy management approaches typically rely on threshold-based mechanisms that do not adequately reflect real-time traffic dynamics, leading

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723105

to inefficient energy usage and unnecessary power wastage during low-demand periods. As energy costs and environmental sustainability concerns continue to grow, there is a critical need for more intelligent and adaptive solutions that can manage energy consumption without compromising the network's availability.

To address this challenge, a RL-based predictive energy management model is proposed here. This model dynamically learns and predicts traffic patterns and makes real-time decisions regarding base station operational modes such as switching between active and sleep states or adjusting transmission power levels. The RL approach enables the system to learn optimal control policies over time through interaction with the network environment, guided by an energy-minimizing reward function. The Markov decision process of the RL contains the state spaces, action space, reward, transition and discount factor as shown in Figure 1.

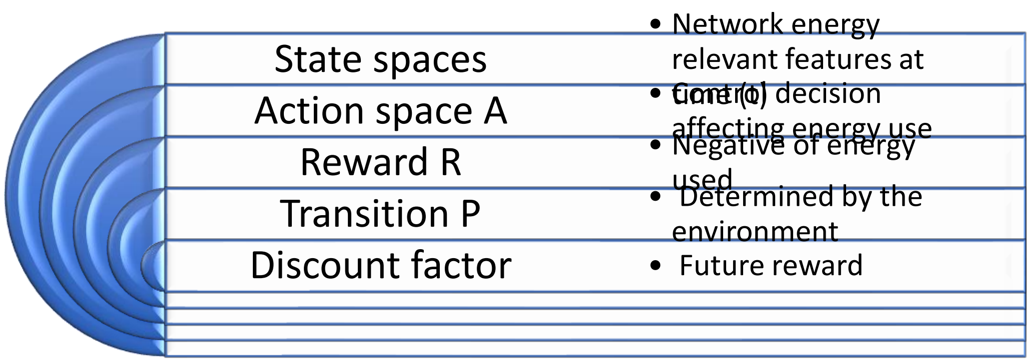

Figure 1: Markov decision process of the RL

RL in energy management operates as a Markov Decision Process, where the agent interacts with the network environment over discrete time steps. At each time t, the agent observes a state comprising energy-relevant features like current traffic load or operating status and selects an action from the action space, such as putting the base station in sleep or active mode. After executing the action, the environment returns a reward, which in this context is the negative of the energy consumed, encouraging the agent to minimize energy usage. The system then transitions to a new state, governed by the environment's behavior (changing traffic), and this transition is considered stochastic, forming the basis of transition probability (P). The agent's goal is to learn a policy that maximizes the expected cumulative reward over time, considering not only immediate energy savings but also future consequences, which are weighted using the discount factor (gamma). Through repeated interaction, the agent learns an optimal energy control strategy that balances present actions with future energy efficiency.

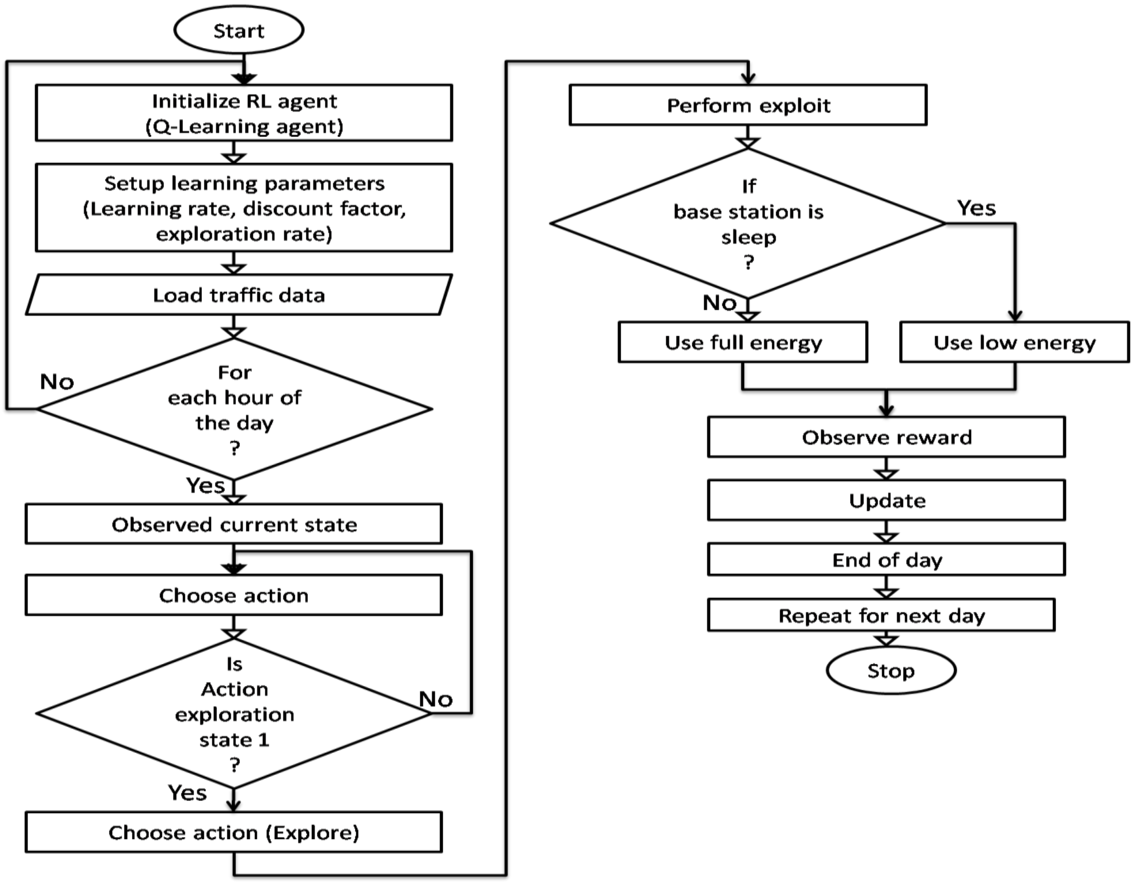

Figure 2: Flowchart of the RL algorithm

Q-learning Agent

Q-learning is a model-free reinforcement learning algorithm used to learn the optimal action-selection policy in a Markov Decision Process (MDP). It does not require a model of the environment and learns directly from interactions. Equation 1 presents the update.

$$Q(s,a) \leftarrow Q(s,a) + \alpha\,[r + y.\max Q(s',a') - Q(s,a) \qquad (1)$$

Where:

- $Q(s,a)$: Current estimate of the value of taking action a in state s
- α: Learning rate (0.1) how quickly the model learns
- r: Immediate reward received (negative energy used)
- γ : Discount factor (0.9) how much future rewards are valued
- s′: Next state
- $\max Q(s',a')$ Best future value possible from next state

Integration of the RL on the 5G Base Station for Energy Management

In this study, an RL-based control mechanism is embedded into the base station's energy management layer using python programming language. The RL agent interacts continuously with the base station environment, observing traffic and power usage every time of day. Based on these observations, it selects optimal actions such as adjusting transmission power, activating/deactivating antenna arrays, scheduling user access dynamically, or enabling low-power modes during idle periods. Figure 3 presents the flow chart of the integrated 5G network with RL.

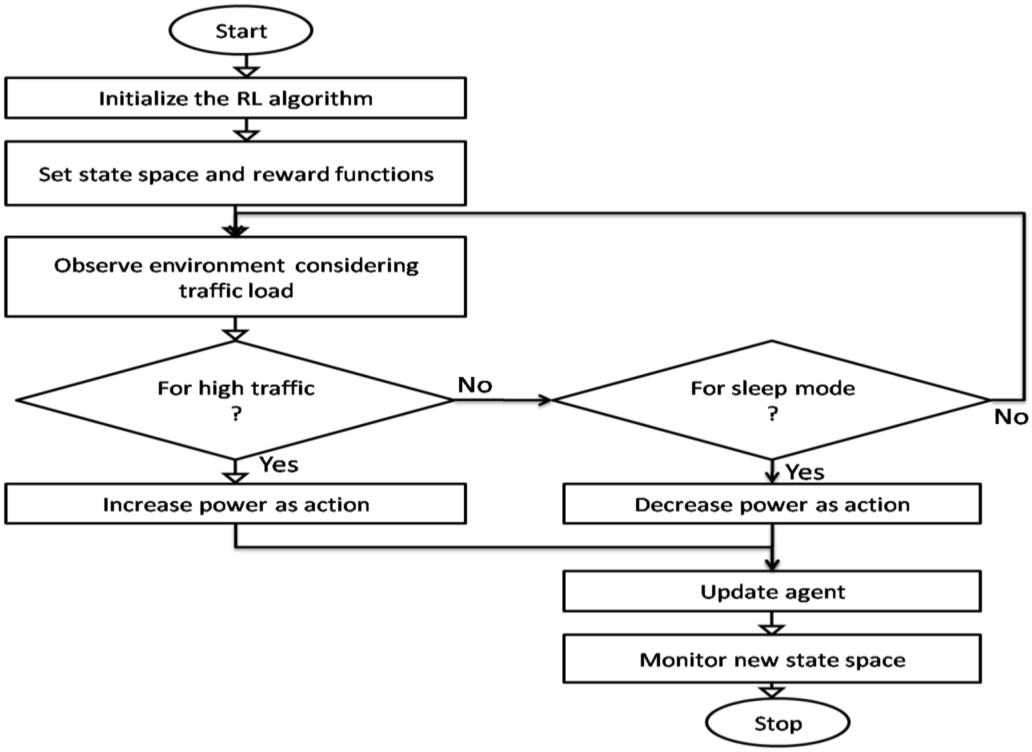

Figure 3: Flowchart of the RL based integration on the 5G network

The updated agent learns optimal policy through continuous reward feedback that reflects goals like minimizing total energy consumption during sleep mode. In the RL, the Markov decision upon initialization of the RL defines states, actions, transition probabilities, and reward functions, and the agent is Q–learning algorithm. This integration allows the base station to autonomously and intelligently adjust its operational parameters in real-time, adapting to fluctuating traffic demands and environmental conditions.

System Implementation

The implementation phase involves translating the proposed reinforcement learning-based energy management framework into an operational model using Python programming. Python is chosen due to its vast ecosystem of libraries for machine learning, reinforcement learning, and scientific computing. The implementation step includes defining the environment using actual simulation parameters collected from the testbed. The Q-leaning agent was used as the RL algorithm to observe the state space of the network environment and then select appropriate action to manage energy. The RL agent is trained over several episodes, with each episode simulating a time period. The agent receives feedback in the form of a reward signal based on the energy saved. During training, performance metrics such as total energy consumption, and convergence of the learning algorithm are logged. Matplotlib library was used to visualize training progress and decision outcomes. The results obtained were evaluated with power consumed per day as the major metric as dynamic load conditions.

III. RESULT OF RL AGENT TRAINING

This section presents result of the RL training process to learn the dynamic environmental behavior of the base station during load and energy consumption. Figure 4 presents the daily energy consumption learning process of the RL agent.

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723105

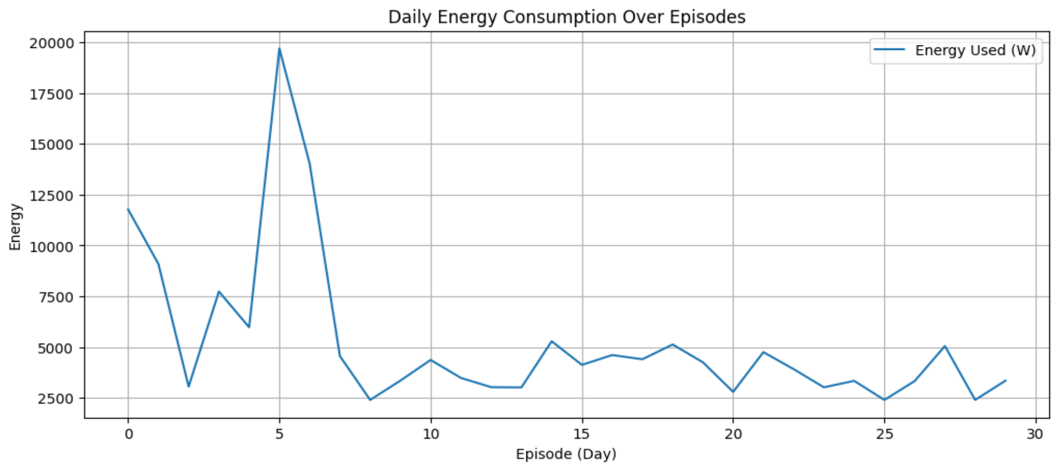

Figure 4: Daily Energy Consumption over Episodes

This graph in Figure 4 illustrates how the total energy consumed by the 5G base station evolves over the course of 30 training episodes (days). Initially, energy consumption is higher due to the RL agent's exploration of different actions without sufficient experience. However, as training progresses, the agent gradually learns more efficient energy-saving policies, leading to a noticeable reduction or stabilization in daily energy use. This trend demonstrates the effectiveness of the RL approach in optimizing operational energy costs under dynamic traffic conditions. Figure 5 presents the cumulative reward per episode by the RL agent.

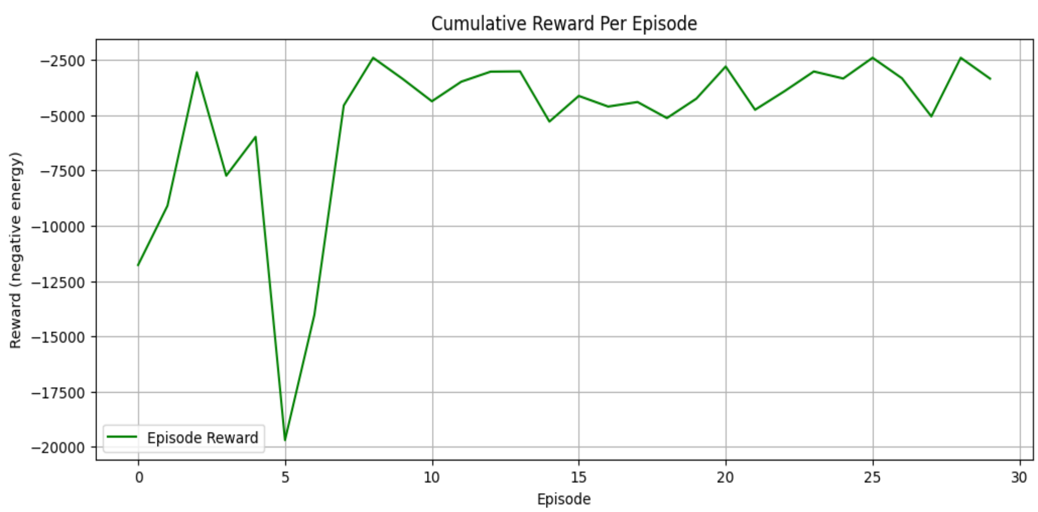

Figure 5: Cumulative reward of the RL agent

The reward graph in Figure 5 shows the total reward collected each day, where the reward is defined as the negative of energy consumed. Hence, a higher reward corresponds to better energy efficiency. The increasing pattern across episodes signifies that the RL agent is learning to minimize energy consumption by selecting actions that lead to lower power use without being explicitly told what the best action is. This progression reflects successful reinforcement learning, where the agent improves its policy through trial, error, and feedback from the environment.

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723105

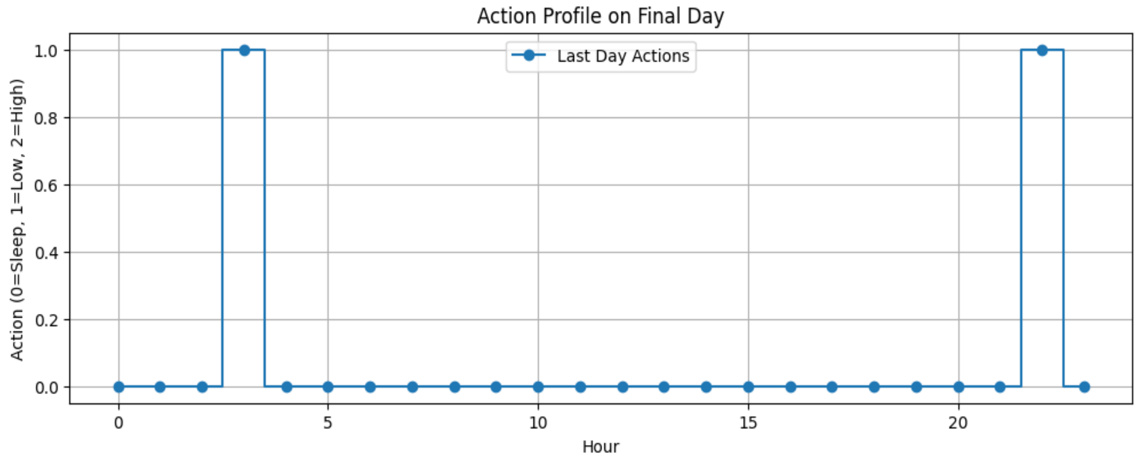

Figure 6: Action Profile on Final Day

This Figure presents the hourly energy management actions taken by the RL agent on the last day of training. The action choices reflect intelligent scheduling, where the agent tends to select sleep or low-power modes during low-traffic periods and switches to high-power mode during high-demand hours. This dynamic and context-aware behavior confirms that the agent has effectively learned to balance energy savings with network demand, showcasing its adaptive decision-making capabilities after full training.

Simulation of the RL on the 5G base station

This section discusses the simulation results obtained from the implementation of the RL-based predictive energy management model on a 5G base station. The simulation was conducted using traffic and environmental data, including hourly user load and temperature trends, collected from MTN's 5G network operations. The RL agent was trained using a Q-learning architecture, designed to learn optimal energy control policies by minimizing energy consumption without compromising service availability. The simulation spans a 30-day period, with hourly variations reflecting realistic network behavior. Figure 7 presents the hourly energy consumed on the site.

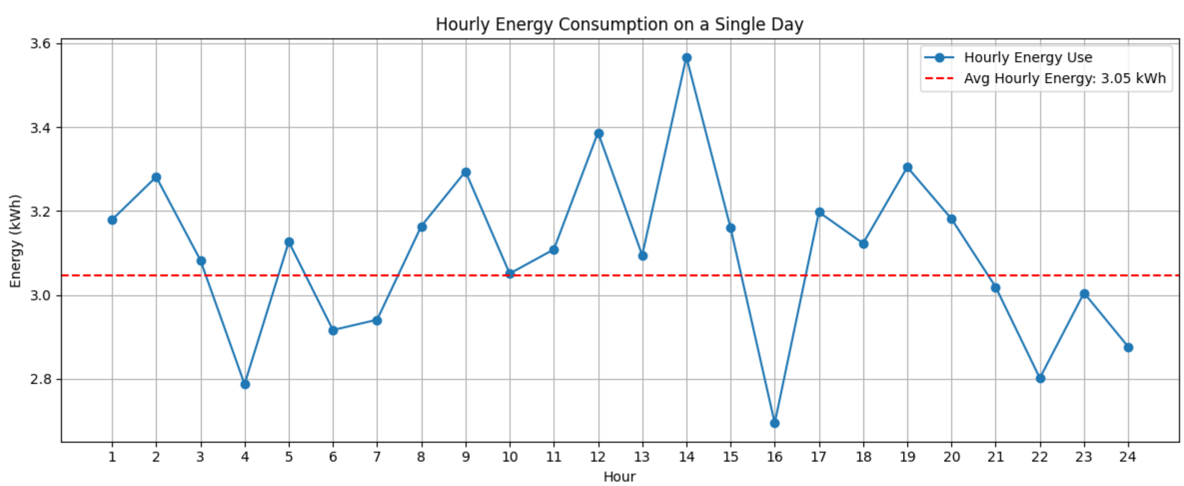

Figure 7: Hourly energy consumption on the network

The graph in Figure 7 presents the "Hourly Energy Consumption on Final Day," and shows the amount of energy used each hour during the final episode (day) of training. Each bar corresponds to the energy consumed based on the RL agent's chosen action: sleep, low-power, or high-power mode. The variation in bar heights indicates the agent's dynamic response to hourly changes in traffic demand activating low or high modes during busy hours and switching to sleep mode during low-demand periods. This hourly analysis confirms that the RL agent not only learns over time but also makes context-aware decisions within each day to optimize energy use. From the results, it was observed that while the energy consumed changes every day, the average hourly energy consumption by the site is 3.05kWh. Comparing the characterized which consumed 4.55686Kw/h of energy per day, it was observed that with adaptation of the base station behavior with RL was able to optimize energy consumed. The percentage reduction in energy consumed is 33.05% with RL. In the Figure 8, the daily energy consumed with RL was simulated for 30 day. The reason was to determine the overall average energy used by the site in a mo nth with RL integrated into it.

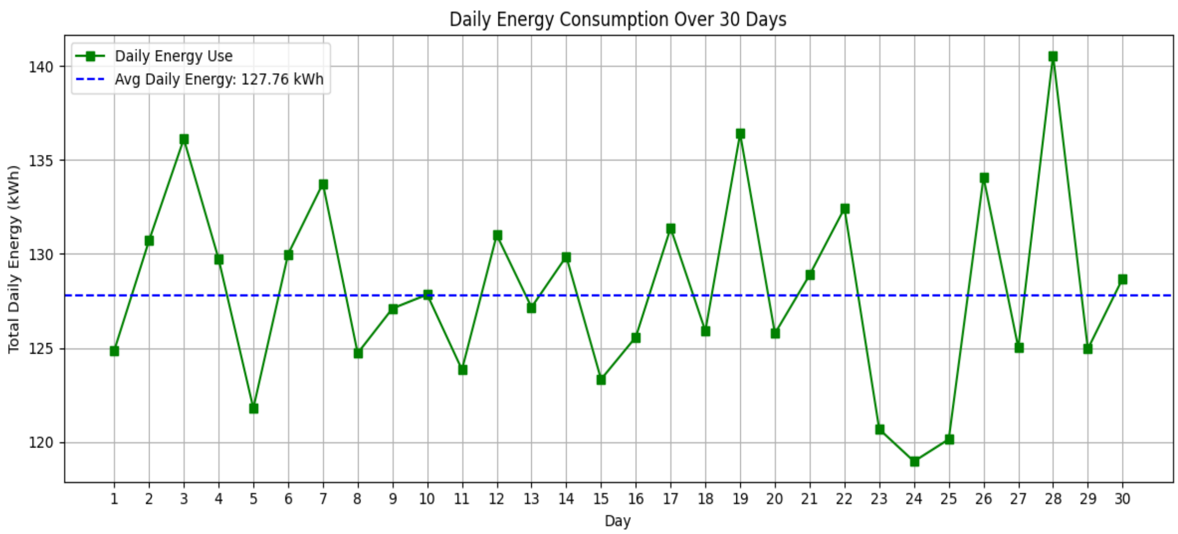

Figure 8: Daily energy consumed over 30 days in the month

The "daily energy consumption over 30 days presents a trend of total energy consumed by the 5G base station across a 30-day simulation period. This graph in Figure 8 reflects the learning progression of the RL agent. Initially, energy consumption was relatively high due to random or suboptimal action selection during the exploration phase. However, as the agent gains experience, it begins to apply more efficient energy control strategies, leading to more consistent and reduced energy usage. The average energy consumed everyday is 127.7kWh as against 149.7746kWh. This trend demonstrates the agent's ability to adapt its policy over time for long-term energy savings. The percentage reduction in energy consumed is 14.74% with RL.

Result of System Integration of the RL-based Predictive Energy Management Model

The system integration section discussed how the RL-based predictive energy management model deployed 5G base station was able to management energy consumed on the network. From the simulation result, 14.74% improvement was recorded within a month as the percentage reduction in daily energy saving. 33.05% reduction was recorded as improvement in hourly energy consumed for the day. The result of the system integration was reported in Table 1.

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723105

Table 1: Result of system integration of RL on the base station for 5G network

Time	Hourly Energy_Consumption_kWh in 5 days					Daily energy Energy_Consumption_kW/day	Normalized Traffic Load
	1	2	3	4	5		
0:00	1.96833	2.10223	2.256215	1.345695	2.216045	115.0754	0.54
1:00	2.93241	2.65122	1.64697	2.70478	2.055365	135.2394	0.82
2:00	2.564185	1.67375	2.95919	2.524015	1.37917	141.9408	0.88
3:00	2.34325	2.202655	2.63783	2.55749	1.519765	149.4011	0.95
4:00	1.600105	2.32986	2.912325	2.631135	1.39256	134.4209	0.81
5:00	1.600105	1.41934	2.83868	1.466205	2.403505	134.4209	0.81
6:00	1.439425	2.35664	2.336555	1.94155	1.867905	124.3006	0.69
7:00	2.791815	1.626885	2.87885	1.533155	2.189265	117.3774	0.59
8:00	2.34325	1.44612	1.48629	2.78512	2.858765	132.4173	0.77
9:00	4.76684	5.255575	3.742505	4.599465	3.849625	141.0456	0.87
10:00	3.38767	5.28905	3.44123	4.010305	4.170985	132.4685	0.78
11:00	5.295745	4.974385	4.00361	3.474705	4.867265	132.4429	0.78
12:00	5.02125	3.956745	4.130815	3.970135	3.809455	138.4793	0.84
13:00	3.77598	3.541655	3.889795	4.00361	3.501485	94.52776	0.35
14:00	3.715725	4.719975	5.014555	4.813705	3.929965	96.13065	0.38
15:00	3.715725	4.23124	4.063865	4.626245	3.66886	131.6244	0.76
16:00	3.956745	3.595215	3.90988	5.12837	5.215405	127.7792	0.72
17:00	4.398615	4.345055	4.438785	4.29819	4.96769	139.0932	0.85
18:00	2.06206	1.399255	1.573325	1.53985	2.39681	128.6744	0.73
19:00	1.827735	2.858765	2.684695	2.53071	2.79851	124.3773	0.68
20:00	2.363335	1.774175	1.466205	2.61105	2.684695	123.3371	0.66

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723105

21:00	1.573325	2.45037	2.992665	2.2763	1.653665	108.9111	0.48
22:00	1.827735	1.86121	2.631135	2.631135	2.831985	136.9958	0.83
23:00	1.95494	2.20935	1.67375	2.162485	2.242825	124.2665	0.68
AVG	2.884407	2.927924	2.983741	3.006925	2.936293	127.6978	0.76

The results presented in Table 1 show the hourly energy consumption (in kWh) across five different test days after integrating the RL-based predictive energy management model into the 5G base station, alongside the average daily energy consumption and the normalized traffic load. The RL agent was tasked with dynamically adjusting the operational mode of the base station switching between active, low-power, and sleep modes based on real-time traffic conditions in order to minimize energy use.

From the table, the average hourly energy consumption over the five days ranges between 2.884 kWh and 3.007 kWh, with an average hourly cost for the five days as 2.947864kWh, indicating a relatively stable yet slightly varying power usage pattern across time. This variation aligns with minor fluctuations in traffic demands which the RL agent uses as part of its state input. Also, none of the hourly values spike excessively, suggesting the RL-based predictive energy management model has successfully avoided energy waste by making context-aware decisions in each time slot. The average daily energy consumption across these days is 127.6978 kWh, which represents a significant reduction compared to the baseline system's daily consumption (often over 149 kWh, as noted in earlier calculations at the characterized testbed). This reflects the RL model’s learned capability to operate efficiently even under moderate traffic loads.

Cost analysis with RL-based predictive energy management model

Since the unit cost of electricity is ₦209.5 for1kWh of electricity, the daily average energy used is 26752.6891kWh. The cost of energy used every day by the site is ₦26752.6891. This value when applied to compute the cost of running the site for 30day (One Month) is ₦802580.673.

IV. CONCLUSION

The cost of power in Nigeria has continued to dominate the new since the emergence of the current administration in 2023. This high cost has continued to impact economically on all telecommunication stakeholders and requires urgent attention. Traditional base station lack effective energy management schemes and this has continued to present economic burden to the network administrators and has also directly affects users as they have to pay more. The light at the tunnel is the proposed and designed RL-based predictive energy management model which allows the site to adaptive to load changes and save energy. In carrying out this task, the first step was investigation of the site to measure hourly, daily and month energy consumed by the site. Cost analysis was also carried out to determine the economic implications. Findings from investigation revealed the cost of daily energy used by the site is ₦31377.7787 and also the monthly running cost is ₦972711.1397. To reduce the economic burden, RL-based predictive energy management model was proposed as an adaptive learning algorithm which utilized the learning agent to adjust the action of the site based on dynamic state space. Python programming language was used to train the model and also integration on the 5G network. Results obtained revealed that with RL, the daily energy saved when compared with characterized is ₦4625.0896. The monthly energy saved amount to ₦170130.4667 with RL-based predictive energy management model, while every years, ₦2041565.6004 is the amount saved. The reason was because of the ability of the RL to adjust the action of RL based on the different state space of the site to save energy.

DOI: https://doi.org/10.64388/IREV10I3-1723105

REFERENCES

[1] Ezzeddine, Z., Khalil, A., Zeddini, B., & Ouslimani, H. H. (2024). A Survey on Green Enablers: A Study on the Energy Efficiency of AI-Based 5G Networks. *Sensors, 24*(14), 4609. MDPI

[2] Ichimescu, A., Popescu, N., Popovici, E. C., & Toma, A. (2024). Energy Efficiency for 5G and Beyond 5G: Potential, Limitations, and Future Directions. *Sensors, 24*(22), 7402. MDPI

[3] Ichimescu, A., Popescu, N., Popovici, E. C., & Toma, A. (2024). Energy Efficiency for 5G and Beyond 5G: Potential, Limitations, and Future Directions. *Sensors, 24*(22), 7402. MDPI

[4] Kaur, P., Garg, R. & Kukreja, V. Energy-efficiency schemes for base stations in 5G heterogeneous networks: a systematic literature review. *Telecommun Syst* 84, 115–151 (2023). Springer

[5] Ma, X., Zhu, Q., Duan, Y., et al. (2022). Optimal configuration of 5G base station energy storage considering sleep mechanism. *Global Energy Interconnection, 5*(1), 66–76.

[6] Premalatha, J., & SahayaAnselin Nisha, A. (2023). Base station energy management in 5G networks using wide range control optimization. *Intelligent Automation & Soft Computing*. Tech Science Press

[7] Wang, Y. (2020). Predictive traffic load forecasting for base station energy management in cellular networks. *IEEE Access, 8*, 74508–74517.

[8] Zhu, Y., Li, K. and Zhang, L. (2025). Base station microgrid energy management in 5G networks - a brief review. In: *Smart Grid and Cyber Security Technologies. The 2024 Intelligent Computing for Sustainable Energy and Environment (ICSEE2024), 13-15 Sep 2024, Suzhou, China. Communications in Computer and Information Science*. Springer, Singapore ISBN 978-981-96-0224-7. Springer