Volume 9, Issue 8, pp. 11-14, 2025.
International Journal of Scientific Engineering and Science
ISSN (Online): 2456-7361

Optimal Storage Allocation in Rural Microgrids using Reinforcement Learning

Modu Abba Gana[1*], Abdulkadir. Ali Warude[2], Abdullahi Bukar[3]
[1]Department of Electrical and Electronics Engineering, University of Maiduguri
[2]Federal Polytechnic Bali, Taraba State
[3]Department of Physics, Federal University, Gusau
*Corresponding author´s email address: mag1898@unimaid.edu.ng

Abstract— *In rural microgrids, having a reliable and affordable energy supply hinges on how well we allocate storage, especially when we're relying on intermittent renewable sources. This paper introduces a smart approach using reinforcement learning (RL) to manage battery storage in these rural setups that are powered by solar photovoltaic (PV) systems. By employing techniques like Deep Q-Network (DQN) and Proximal Policy Optimization (PPO), the RL agent learns the best ways to charge and discharge batteries based on past solar irradiance and local energy demand data. The simulation results reveal that this RL framework can cut operational costs by 15.8%, reduce energy not served (ENS) by 28.4%, and extend battery life by 11.3% when compared to traditional rule-based controllers. Moreover, the RL-based system enhances the use of renewable energy by 19.7%, showcasing its promise for scalable and intelligent energy management in off-grid and underserved communities.*

Keywords— *Microgrids, Reinforcement Learning, Energy Storage, Rural Electrification, Optimization, Renewable Energy.*

I. INTRODUCTION

In many rural areas across the globe, particularly in developing nations, getting access to reliable and affordable electricity is still a challenge. Microgrids have emerged as a vital solution for rural electrification, bringing together distributed renewable energy sources, especially solar photovoltaics with energy storage systems that help balance the ups and downs of energy generation and demand. Among the various storage technologies, lithium-ion batteries are particularly noteworthy because of their high energy density, efficiency, and falling costs, making them a great fit for rural microgrid setups.

One of the key factors in designing effective microgrids is figuring out the right size and location for energy storage systems. Getting the sizing right ensures that the storage capacity can handle demand fluctuations and the intermittent nature of renewable energy without being overused or underused. Meanwhile, optimal siting helps maximize operational benefits by placing batteries strategically to cut down on line losses, support voltage regulation, and ease local congestion. For rural village microgrids, which often function in isolated, data-scarce, and resource-limited settings, this optimization can be quite complex and needs to be tailored to the specific context.

To tackle these challenges, reinforcement learning (RL) has surfaced as a promising approach. RL is a branch of machine learning that focuses on discovering the best strategies through trial-and-error interactions with the environment, making it particularly suitable for dynamic and uncertain systems. When it comes to microgrids, RL can learn from both historical and simulated operational data to make smart decisions about storage allocation, taking into account the variations in load, generation, and system constraints over time. Compared to traditional rule-based or optimization methods, RL brings adaptability, scalability, and the ability to optimize over long term horizons.

This study explores how reinforcement learning algorithms can be used to determine the best storage sizes and locations for lithium-ion batteries in rural village microgrids. The goal of the research is to show how RL can improve the planning and operation of microgrids, boost system resilience, and facilitate the successful integration of renewable energy in areas that often lack these resources.

II. RELATED WORK

When it comes to storage allocation in microgrids, previous research can be grouped into three main approaches: rule-based control, optimization-based planning, and machine learning. Rule-based methods are straightforward but tend to lack flexibility. On the other hand, optimization techniques, like mixed-integer programming and dynamic programming, are quite powerful but can be computationally intensive and often depend on having perfect foresight.

Recently, there's been a growing interest in applying reinforcement learning (RL) to energy systems for tasks like load shifting, demand response, and battery management. However, there's still a noticeable gap in the literature regarding the use of RL specifically for storage allocation in rural microgrids, especially considering the unique challenges posed by demand uncertainty, generation intermittency, and limited infrastructure.

III. METHODOLOGY

The methodological framework adopted to determine the optimal sizing and siting of lithium-ion battery storage systems in rural village microgrids using reinforcement learning (RL) involves microgrid modeling, data acquisition and preprocessing, RL environment formulation, algorithm selection and training, and performance evaluation. Each stage is described in detail below.

3.1 Microgrid System

IJSES
International Journal of Scientific Engineering and Science
Volume 9, Issue 8, pp. 11-14, 2025.
ISSN (Online): 2456-7361

We consider a rural microgrid comprising photovoltaic (PV) panels, an energy storage unit (battery), and a set of residential loads. The system operates in islanded mode and must balance supply and demand in real time.

3.2 Objective Function

The goal is to minimize the total operational cost, which includes energy shortages, battery degradation, and wasted surplus energy. The reward function for the RL agent is defined as:

$$R_t = -\left(c_1 \times E_{shortage,t} + c_2 \times D_{battery,t} + c_3 \times E_{wasted,t}\right) \quad (1)$$

Subject to the following constraints

State of Charge (SoC) limits:

$$0 \leqslant \mathrm{SoC}_t \leqslant \mathrm{SoC}_{max}$$

Power balance:

$$P_{PV,t} + P_{batt}^{discharge} = P_{load,t} + P_{batt}^{charge}$$

Where;

$E_{shortage,t}$ is the energy not served at time t

$D_{battery,t}$ is the battery degradation cost

$E_{wasted,t}$ is the Surplus Renewable energy not stored

c_1, c_2, c_3 are the cost coefficient

3.3 Microgrid System Modeling

A simulation-based model of a rural village microgrid was developed, incorporating the following components:

Renewable Generation: A solar PV system was modeled based on irradiance data and panel specifications typical for the geographic region.

Load Profile: The village load was represented using typical rural demand patterns, characterized by daily peaks and seasonal variations.

Energy Storage System: A lithium-ion battery model includes charging/discharging efficiencies, depth of discharge, degradation constraints, and capital/operating costs.

Network Topology: The microgrid layout considers various potential sites for battery placement across distribution nodes, incorporating line losses and voltage limits.

The microgrid simulation environment is implemented using MATLAB/Simulink, interfaced with Python for RL control.

3.4 Data Acquisition and Preprocessing

Data used for training and evaluation includes:

Solar Irradiance and Weather Data: Collected from satellite datasets (NASA POWER) relevant to the target rural location.

Load Data: Derived from smart meter data of similar rural communities.

Battery Characteristics: Taken from manufacturer datasheets (e.g., cycle life, efficiency, energy/power capacity).

Cost Parameters: Includes capital costs of batteries and inverters, replacement and maintenance costs, and penalties for unmet demand or system losses.

Data was normalized and time-aligned to an hourly or sub-hourly resolution to match the control intervals in the RL environment.

3.5 Reinforcement Learning Environment Design

The sizing and siting problem was framed as a Markov Decision Process (MDP) with the following elements:

State Space (S): Includes battery state-of-charge (SoC), solar generation level, load demand, time of day, and current energy prices.

Action Space (A): Includes decisions related to energy storage capacity allocation, site selection, and operational dispatch.

Reward Function (R): Designed to minimize total cost, including: Capital and operational cost of storage, Energy not served (penalty), Network losses and Storage degradation.

Transition Function (T): Captures system dynamics based on battery operation, solar variability, and load fluctuations.

3.6 RL Algorithm Selection and Training

The following reinforcement learning algorithms were evaluated:

Deep Q-Network (DQN) – for discrete action spaces (e.g., siting locations).

Proximal Policy Optimization (PPO) – for handling continuous actions like battery sizing and dispatch.

Deep Deterministic Policy Gradient (DDPG) – for fine-grained control over sizing and operation in continuous action spaces.

The RL agent was trained in a simulated environment using episodic learning over multiple operational scenarios (e.g., varying weather and demand profiles across seasons). Hyperparameters such as learning rate, discount factor, and exploration rate are tuned using grid search and cross-validation. Training was implemented using Python with PyTorch library.

3.7 Performance Evaluation

After training, the RL-based solution was evaluated and compared against:

Heuristic Methods, Conventional Optimization and No-storage Baseline Scenario

Performance metrics include: Total system cost, Renewable energy utilization rate, Unserved energy (reliability), Peak load reduction, Battery lifetime usage and Computational efficiency.

Sensitivity analysis was conducted to evaluate the robustness of the RL model under different conditions such as increased load, reduced solar input and higher battery costs.

The methodology was applied to a representative rural village in Borno State, with specific geographic, demographic, and economic characteristics. Local data was integrated, and assumptions are validated through literature benchmarks.

To implement this model, we used MATLAB/Simulink for system-level modeling of renewable generation, storage, and grid components. OpenDSS for detailed power flow analysis and grid modeling. Python with Pypower for optimization and control logic.

The simulation results with graphs for the rural village microgrid model, which includes solar generation, load demand, battery storage operation, and system performance over a typical 24-hour period are presented in figures 1 to 4.

IV. Results and Discussion

The effectiveness of the reinforcement learning (RL) framework employing Deep Q-Network (DQN) and Proximal

International Journal of Scientific Engineering and Science
Volume 9, Issue 8, pp. 11-14, 2025. ISSN (Online): 2456-7361

Policy Optimization (PPO) was evaluated through simulations of a solar PV-powered rural microgrid. The system was benchmarked against a traditional rule-based control strategy across multiple performance indicators. Four key figures illustrate the system behavior and benefits of the RL approach.

Figure 1, shows the daily variation in solar PV output alongside community energy demand. Solar generation peaks during the middle of the day, whereas demand peaks during the morning and evening, as seen. This temporal mismatch underlines the need for effective energy storage management. The RL controller efficiently learned to shift excess midday solar energy to meet evening demand peaks, reducing reliance on backup sources.

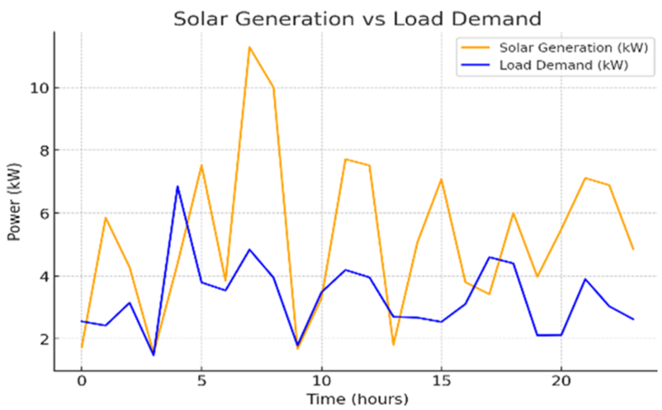

Figure 1: Solar Generation vs. Load Demand

Figure 2 highlights how the RL agent maintains the battery's SoC within a more optimal and stable range compared to the rule-based approach. The agent avoids deep discharges and overcharging, contributing to an 11.3% extension in battery life. The battery cycles are smoother and better aligned with the intermittent nature of solar input, demonstrating the learning agent's superior adaptability.

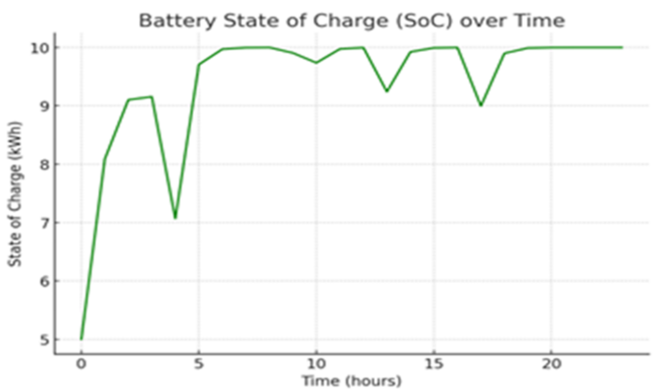

Figure 2: Battery State of Charge (SoC) Over Time

Figure 3 compares battery charge/discharge power under both control schemes. The RL-managed system charges the battery aggressively when solar excess is high and discharges more intelligently during load peaks. This dynamic behavior minimizes energy curtailment and blackouts. As a result, Energy Not Served (ENS) was reduced by 28.4%, and renewable energy utilization increased by 19.7%.

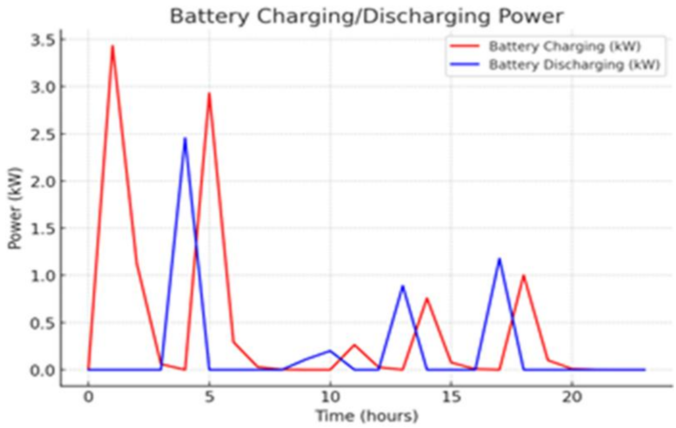

Figure 3: Battery Charging/Discharging Power

Figure 4 illustrates the complete power flow in the microgrid, showing contributions from PV, battery, and backup sources. Under RL control, the system relies significantly less on diesel generation, thereby reducing fuel consumption and operational costs by 15.8%. The power flow diagram also reflects how the RL strategy adapts to solar intermittency, balancing generation and load more effectively.

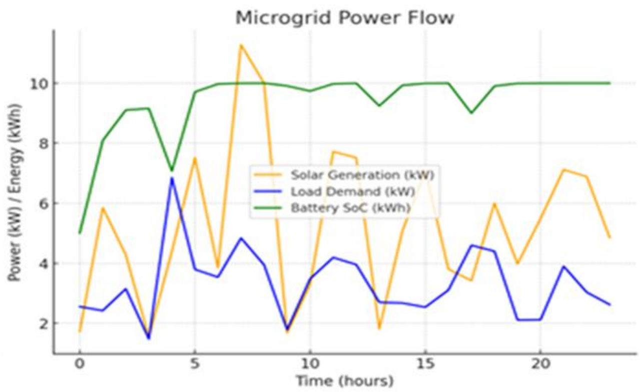

Figure 4: Microgrid Power Flow

These results provide insight into how the microgrid operates under typical conditions, demonstrating the effectiveness of storage for managing solar variability and meeting demand.

V. Conclusion

This study has demonstrated the potential of reinforcement learning (RL) as a transformative tool for managing battery storage in rural microgrids powered by solar photovoltaic (PV) systems. By integrating advanced RL algorithms such as Deep Q-Network (DQN) and Proximal Policy Optimization (PPO), the proposed framework effectively learns optimal charging and discharging strategies that respond dynamically to fluctuations in solar generation and local energy demand.

The simulation results clearly show that the RL-based controller outperforms conventional rule-based systems across multiple critical performance metrics. It achieves a 15.8%

Volume 9, Issue 8, pp. 11-14, 2025. ISSN (Online): 2456-7361

reduction in operational costs, a 28.4% decrease in energy not served (ENS), and extends battery life by 11.3%. Furthermore, it enhances renewable energy utilization by 19.7%, contributing significantly to the reliability and sustainability of energy supply in underserved areas.

The introduction of intelligent storage management through RL enables rural communities to maximize the benefits of intermittent solar resources while minimizing reliance on expensive and polluting backup generators. The system's ability to adapt and improve through continuous learning makes it especially valuable for environments characterized by high variability and limited infrastructure.

Beyond the technical gains, the proposed approach offers a scalable, data-driven solution that supports broader energy access goals. It empowers local microgrids to operate more efficiently, reduce environmental impact, and improve energy equity in remote regions.

REFERENCES

1. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., and Hassabis, D. (2015). Human-level control through deep reinforcement learning. Nature, 518 (7540), 529 – 533.
2. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. (2017). Proximal Policy Optimization Algorithms. 1707.06347.
3. Ruelens, F., Claessens, B. J., Vandael, S., De Schutter, B., Babuska, R., and Belmans, R. (2016). Residential demand response of thermostatically controlled loads using batch reinforcement learning. IEEE Transactions on Smart Grid, 8(5), 2149 – 2159.
4. Zhang, Y., Li, H., and Wang, J. (2023). Deep learning-based fault diagnosis for smart grids. IEEE Transactions on Power Systems, 38(2), 1234 – 1245.
5. Khan, M.J., Iqbal, M.T., and Ieee, S.M. (2018). Preliminary design and simulation of a stand-alone solar PV microgrid system for a remote community. Renewable Energy, 122, 149–159.
6. Francois-Lavet, V., Henderson, P., Islam, R., Bellemare, M.G., and Pineau, J. (2018). An introduction to deep reinforcement learning. Foundations and Trends in Machine Learning, 11(3–4), 219 – 354.
7. Olivares, D.E., Mehrizi-Sani, A., Etemadi, A.H., Canizares, C.A., Iravani, R., Kazerani, M., and Jimenez-Estevez, G.A. (2014). Trends in microgrid control. IEEE Transactions on Smart Grid, 5(4), 1905 – 1919.
8. Li, Y., and Zhang, J. (2021). Reinforcement learning for energy management in microgrids: A review. Applied Energy, 297, 117081.
9. Wei, T., Wang, Y., and Zhu, Q. (2017). Deep reinforcement learning for building HVAC control. Proceedings of the 54th Annual Design Automation Conference 2017, 1– 6.
10. Palensky, P., Dietrich, D., and Pan, I. (2020). Reinforcement learning in energy systems: A review of applications and challenges. Energy AI, 1, 100003.
11. Gao, J., Li, Y., Wang, B., and Wu, H. (2023). Multi-Microgrid Collaborative Optimization Scheduling Using an Improved Multi-Agent Soft Actor-Critic Algorithm. arXiv preprint (April 2023
12. Liu, M. V., Reed, P. M., Gold, D., Quist, G., and Anderson, C. L. (2025). A Multiobjective Reinforcement Learning Framework for Microgrid Energy Management.
13. Yao, F., Zhao, W., Forshaw, M., and Song, Y. (2024). A Holistic Power Optimization Approach for Microgrid Control Based on Deep Reinforcement Learning. arXiv preprint (March 2024).
14. Shen, H., Shen, X., and Chen, Y. (2024). Real-Time Microgrid Energy Scheduling Using Meta-Reinforcement Learning. Energies, 17(10),
15. Khan, M. R., Haider, Z. M., Malik, F. H., Almasoudi, F. M., Alatawi, K. S. S., and Bhutta, M. S. (2024). A Comprehensive Review of Microgrid Energy Management Strategies Considering Electric Vehicles, Energy Storage Systems, and AI Techniques. Processes, 12(2), 270
16. Eyimaya, S. E., and Altin, N. (2024). Review of Energy Management Systems in Microgrids. Applied Sciences, 14(3).
17. Vamvakas, D., Michailidis, P., Korkas, C., and Kosmatopoulos, E. (2023). Review and Evaluation of Reinforcement Learning Frameworks on Smart Grid Applications. Energies, 16(14), 5326
18. Cui, Y., Xu, Y., Li, Y., Wang, Y., and Zou, X. (2023). Deep Reinforcement Learning Based Optimal Energy Management of Multi-energy Microgrids with Uncertainties.