

Optimizing industrial energy consumption through artificial intelligence: a study on cost reduction and sustainability

Sri Harsha Giri^{1*}, Dr. Nilam Nimraj Ghuge²

¹ Research Scholar, Department of Electrical, University Of Technology, Jaipur, Rajasthan, India

harshaguys1988@gmail.com

² Professor, Department of Electrical Engineering, University Of Technology, Jaipur, Rajasthan, India

Abstract : Effective power consumption control is essential for the steel industry, particularly for direct reduction (DR) units, in order to cut operating costs and limit the sector's impact on the environment. The reason for this is because the steel industry is one of the industrial sectors that employs the greatest amount of energy. Through the use of three cutting-edge techniques, namely Response Surface Methodology (RSM), Radial Basis Function (RBF) neural networks, and Multilayer Perceptron (MLP) neural networks, this study investigates the modeling and optimization of power consumption in direct reduction units. We made use of operational data from industrial processes in order to develop models that were capable of estimating the amount of electricity that was used under a variety of process conditions. The comparison demonstrates that artificial intelligence-based models, particularly RBF and MLP networks, deliver greater prediction accuracy when compared to traditional statistical approaches by a significant margin. It was necessary to apply the optimized model once again in order to determine the optimal operating parameters that would minimize the amount of power used without compromising the efficiency of production. Using optimization strategies that are driven by artificial intelligence might potentially result in significant cost reductions and an increase in energy efficiency, according to the research. The purpose of this project is to combine advanced predictive modeling with industrial energy management systems in order to contribute to sustainable production techniques and to promote the move toward energy-efficient steel manufacturing that is compatible with global sustainability goals.

Keywords: Artificial Intelligence, Electricity Consumption Optimization, Direct Reduction Units, Steel Industry, Machine Learning, Energy Efficiency, Sustainable Manufacturing, Multilayer Perceptron (MLP), Radial Basis Function (RBF), Response Surface Methodology (RSM)

INTRODUCTION

One of the most energy-intensive businesses, the steel industry's direct reduction (DR) machines use a considerable amount of the world's electricity. Direct reduction units have high and variable energy needs, making power modeling and optimization challenging. This issue

must be addressed to boost energy efficiency, lower operational costs, and reduce environmental impact. Power consumption must be managed effectively to make steel production in line with global sustainability standards. Accurate power usage forecasting and optimization enable sustainable industrial energy management. Studies have examined many approaches for forecasting and reducing industrial power demand. Using historical data and advanced techniques like time-series regression and gradient descent, one study optimized energy utilization and reduced power use and enhanced sustainability [1]. Deep learning and machine learning models like decision trees and neural networks have improved electricity estimates by up to 15% in the steel sector [2].

Researchers used PCA and regression trees to optimize input data and minimize duplicate variables. Integration improved prediction accuracy by 12% [3]. Industry 4.0's automation, real-time monitoring, and smart manufacturing technologies enable energy optimization. Digital transformation has reduced steel production power usage by 10% [4]. Decision trees, neural networks, and linear regression improve monthly power consumption projections by 20% in industrial applications [5]. Even with little data, AI approaches like neural networks and transfer learning have improved MPE by 8% [6]. In optimising energy distribution systems, graph-based machine learning methods that combine geographical and industrial context enhance forecasting accuracy by 10% [7].

Reviews demonstrate that deep learning and ensemble learning models minimize forecasting errors by 15% or more compared to statistical techniques [8]. Blending Long Short-Term Memory (LSTM) networks with Autoregressive Integrated Moving Average (ARIMA) models reduced errors by 10% [9]. Ensemble learning systems like Random Forest and Gradient Boosting enhance electricity forecasting accuracy by 12% [10]. Advanced control systems, process optimization, waste heat recovery, and predictive modeling have enhanced steel mill energy efficiency by 20%. With 15% power savings, Industry 4.0 and smart energy management technologies are promising. High investment costs and technological complexity prevent its widespread use [11].

Modern AI models like Kolmogorov-Arnold Networks (KAN) exceed multilayer perceptrons (MLP) in interpretability and modeling accuracy. Incorporating the Pearson Correlation Coefficient (PCC) into loss functions improves prediction abilities for large-scale power plants, resulting in $42.17 \pm 0.88\%$ thermal efficiency and reduced turbine heat rates [12]. AI-

powered optimization frameworks in industrial facilities have greatly reduced emissions, operating fuel costs, and efficiency. This shows how AI may help high-emission sectors reach net-zero emissions [13]. Despite these advances, there is little study comparing AI approaches for optimizing power usage in controlled direct reduction units. Additionally, the best predictive approach must be found to estimate power use and discover the ideal operating conditions to decrease power consumption without compromising production efficiency.

This research compares Response Surface Methodology (RSM), Radial Basis Function (RBF), and Multilayer Perceptron (MLP) neural networks for optimizing and forecasting power usage in direct reduction units. This research analyzes model accuracy and robustness to find the optimum industrial approach. We then simulate numerous operational scenarios using the revised model to discover the greatest power consumption-reduction settings without compromising production performance. This study uses advanced artificial intelligence modeling to control steel industry energy and save costs.

OBJECTIVES

1. To create and evaluate neural network models for RSM, RBF, and MLP in order to accurately anticipate power consumption in direct reduction units.
2. To reduce power usage while preserving production efficiency and advancing sustainability objectives by optimizing operational parameters using the best-performing model.

RESEARCH METHODOLOGY

1. Research Design

This study used a quantitative experimental research approach [14] in order to establish the efficacy of energy optimization techniques that are based on Artificial Intelligence (AI) in terms of reducing energy expenditures for industrial processes and increasing sustainability. A comparison is made between three different control methods in the research:

- Logic-based optimization's first algorithm, which was developed by
- PPO Algorithm for Reinforcement Learning (RL) is the second algorithm here.

- The Third Deep Deterministic Policy Gradient (TD3), also known as the Twin Delayed Gradient

Our analysis compares the effectiveness of these algorithms under both the normal and peak pricing structures for electricity.

2. System Architecture

In the course of this investigation, an EMS (Industrial Energy Management System) that includes the following components is discussed:

- Batteries-integrated energy storage system (also known as BESS)
- Establish a connection. The need for load in the industrial sector
- Alterations in the price of electricity

It is the responsibility of the EMS to control the charging and draining of batteries in order

- To reduce demand spikes.
- Ease the burden of the electricity bill.
- Improve the effectiveness of the procedures

3. Data Collection

Information about the hourly industrial load from February 2024 to July 2024. The current hourly rate of electricity, along with its specifics. Specifications about the battery, including its capacity, power constraints, and state of charge limits. Information on fines for peak demand. During the training of the algorithm as well as the simulation, the dataset was used.

4. Battery Scheduling with Peak Shaving

An algorithm for dynamic peak-shaving was used in the following cases:

- If daily maximum load ≥ 200 kW $\rightarrow$ Threshold = 150 kW
- If daily maximum load ≤ 150 kW $\rightarrow$ Threshold = 100 kW
- Otherwise $\rightarrow$ Threshold = 150 kW

When either the load or the electricity price hits their respective thresholds, the battery will drain. The process of recharging batteries takes happen during periods of low interest rates.

- Constraints placed on SOC: SOR must be maintained at a minimum of twenty percent.

5. Algorithm Implementation

Logic-Based Optimization:

- An strategy to control that is outlined in regulations.
- This was accomplished by the use of deterministic decision logic.
- In order to do simulations, the MATLAB-Simulink environment was used.

Reinforcement Learning (RL – PPO):

Proximal Policy Optimization (PPO) is the approach that is used in this process. You may expect a skilled representative to work on lowering your monthly electricity payment to the lowest possible amount.

A few of the goals of the reward function are as follows:

- The heavy grid import should be discouraged. Take steps to reduce demand peaks.
- Encouraging the reduction of expenditures
- It was Python that was used for the training.

TD3 Algorithm (Peak Pricing Scenario):

Twin Delayed Deep Deterministic Policy Gradient (TD3)

Uses:

- Two Q-networks
- Delayed policy updates
- Action smoothing with noise

Compared with PPO under peak-pricing limitations.

6. Simulation Environment

In order to carry out the co-simulation, MATLAB-Simulink was used. A mathematical model coded in Python that was built Performance metrics are evaluated on a monthly basis throughout the year. Algorithms that were tested over a period of six hundred days.

7. Performance Evaluation Metrics

We used the following measures:

1. Monthly Electricity Cost
2. Cost Savings (%)
3. Grid Import Reduction
4. Peak Power Reduction
5. Battery SOC Stability
6. Comparative Performance Between Algorithms

8. Comparative Analysis

Research compares:

- No battery optimization
- Logic-based optimization
- RL (PPO)
- TD3 (Peak pricing)

Savings were calculated as:

Cost Savings=Cost_{without optimization} – Cost_{with optimization}

The computation of percentage improvement was also done.

9. Validation

Validation of algorithms was performed by using real hourly data for a period of six months. For the purpose of comparing performance, bar charts and normalized plots were used.

In accordance with RL,

- ~86% higher savings vs no optimization
- ~20% higher savings vs logic-based optimization

RESULT

Peak Shaving Battery Scheduling

Autonomous charging and discharging schedules for batteries are obtained using the peak shaving algorithm. By following this plan, the EMS may connect with the BESS in an orderly and sophisticated manner. The results of the logic-based approach for automated scheduling with peak shaving [15] are shown in Figure 1 for a day in July. These metrics include Grid import, site load, energy price, battery power, and SOC. This is where the daily peak shaving value is calculated using dynamic peak shaving logic. The dynamic peak-shaving algorithm calculates the cutoff for managing peak power consumption depending on the maximum expected demand for a certain day. The variable 'daily-max-load' captures the maximum expected daily power usage. Afterwards, this number is used to dynamically change the peak shaving threshold. The 150 kW barrier is used when the daily maximum load is 200 kW or above; the 100 kW threshold is used when the daily maximum load is 150 kW or lower. For load forecasts that fall within these ranges, a barrier of 150 kW is maintained. By making optimal use of battery storage, this technique allows for a flexible response to changing load situations while minimizing peak power pricing.




Figure 1. Import into the grid, site load, energy pricing, battery power, and state of charge overnight

It is demonstrated in Figure 1 [16] that the battery drains under two different conditions and charges throughout the day when the cost of electricity is relatively cheap. In the first case, when the load or consumption of the site reaches a specific level, and in the second case, when the energy prices for the day reach their highest point. Furthermore, it ensures that the value of the SOC is larger than twenty percent. To put it another way, when the amount of consumption exceeds the threshold, the algorithm makes efficient use of the battery, which is evidence that it successfully performs demand response. After this, the battery will finish depleting, which will restrict the quantity of grid import to that amount for that hour.

Results of Logic-Based Algorithms

Beginning in February 2024 and continuing through July 2024, the algorithms that were developed were tested and put into effect with the use of monthly hourly data. As can be seen in the bar chart in Figure 2, the logic-based optimization system was able to save money over a period of six months when compared to the alternative of not optimizing the battery at all [17].




Figure 2. Logic-based Optimization Algorithm saves money every month.

When compared to the situation in which the batteries are not optimized, an optimization of the batteries results in considerable savings, as shown by the monthly breakdown. It is during the months of March and February that people are most likely to save anything. On the other hand, when battery optimization is used during the months of May and June, the cost is

negative, which indicates that selling excess energy might be a lucrative business enterprise [18].

Results of the RL Algorithm

Figure 3 illustrates how an algorithm that uses reinforcement learning is able to effectively handle the energy management system. At this point, saving money with optimized batteries (shown by green bars) is always superior than saving money without optimization (represented by blue bars) in each and every month. This graph demonstrates a steady month-over-month decrease in energy expenditures [19]. This is accomplished by proving that the RL Algorithm is more constant and favorably performing than the Optimization Algorithm, which is characterized by its variable performance. A comparison of the performance of the two techniques from February to July revealed that the RL approach consistently produced more cost savings than the Optimization algorithm.

This was the conclusion reached after analyzing the data. During the course of the six months that were investigated, the RL Algorithm led to an increase in average savings of 86%. The significance of this study indicates how successful RL is as a tool for optimization. In Figure 4, a comparison is made between the savings that were achieved via the use of RL and logic-based optimization, which are two different methodologies. The logic-based strategy (yellow) and the RL method (green) both do very well in comparison to one another during the months of February and July. It was estimated that the RL Algorithm had cut expenses by about twenty percent more than logic-based optimization had done by the time the six months were over [20].




Figure 3. Monthly Savings Achieved via the RL Framework.




Figure 4. Comparison of the Two Algorithms' Monthly Savings Results

Peak-Pricing Scheme-Based Economic Optimization

Since the agent must also exchange power with the main grid and ensure that the peak power obtained from the grid for a given month is not too high—otherwise, the overall cost would be greatly increased—the peak pricing scenario presents a far more challenging optimization

challenge. Results produced in the conventional pricing scheme are much higher than those obtained while training the RL agent using PPO. As a result, a separate method known as Twin Delayed Deep Deterministic Policy Gradient (TD3), which is an upgrade to the original DDPG with three major enhancements, was also used to train this scheme. To start, it employs two Q-functions instead of one, which is why it's called "Twin." Then, it updates the Q-functions less often, which is why it's called "Delayed." Lastly, it introduces random action noise to smooth out the actions [21]. We used MATLAB-Simulink and a mathematical model we developed in Python to co-simulate and evaluate PPO and TD3's performance.

Both the TD3 and the PPO agent's normalized spot prices, SOC, and battery power are shown in Figure 5. In comparison to TD3, the PPO agent exhibits less aggressive behavior, as can be shown from these data. At price peaks, like 10 and 40 hours, the two algorithms drain the buffer and at troughs, like 25 hours, they fill it back up. But as soon as the price begins to go up, the TD3 agent appears to switch to discharging. At the 25-second mark, it exhibits a sharp dip before rising again, as the PPO agent deftly tracks price changes. The figures 6(a), (b), (c), and (d) show comparisons of several power metrics, including battery power, energy cost, grid power, and state of charge. In comparison to PPO, these results also show that TD3 is more aggressive. Having said that, there is one noteworthy aspect of this conduct. The TD3 positive peaks are definitely smaller than PPO, as can be shown by examining the grid power in Figure 6(c). Accordingly, TD3 is trying to peak-shave in order to lower the peak penalty that will be applied at the month's conclusion, whilst PPO seems to be trading more in accordance with the current price alone. This might provide light on the reason for TD3's superior performance. Trading at peak penalty/spot price is more difficult to optimize, and training sessions are longer as a result.




Figure 5. The spot price, state of charge, and battery were normalized.




Figure 6. Pbattery, power cost, Pgrid, and SOC comparison between PPO and TD3.

CONCLUSION

The results of this study indicate that the steel industry may be able to maximize the use of electricity in direct reduction units by using modeling strategies that are based on artificial intelligence. When compared to more traditional statistical methods, machine learning models, notably RBF and MLP neural networks, were able to attain a better level of prediction accuracy and made it possible to identify the optimal operating state with more precision. The alterations made to the operating parameters resulted in considerable reductions in the amount of electricity that was used, which in turn led to a reduction in operational expenditures and a positive impact on the environment. It has been shown that industrial energy management systems that make use of predictive analytics powered by artificial intelligence are able to improve decision-making, energy efficiency, and sustainability in the manufacturing sector. According to the findings of the research, the use of advanced modeling tools is of utmost significance in order to accomplish steel manufacture that is not only economical but also environmentally responsible.

References

1. Opoku, R., Obeng, G. Y., Osei, L. K., & Kizito, J. P. (2022). Optimization of industrial energy consumption for sustainability using time-series regression and gradient descent algorithm based on historical electricity consumption data. *Sustainable Analytics and Modeling*, 2, 100004. <https://doi.org/10.1016/j.samod.2022.100004>
2. Kerdprasop, K., Kerdprasop, N., & Chuaybamroong, P. (2023). Deep learning and machine learning models to predict energy consumption in steel industry. *International Journal of Machine Learning*, 13(4), 142–145. <https://doi.org/10.18178/ijml.2023.13.4.1142>
3. Septiani, M., Hoendarto, G., & Tjen, J. (2025). Energy consumption prediction in the steel industry using principal component analysis and regression tree methods. *Social Science and Humanities Journal*. <https://doi.org/10.18535/sshj.v9i01.1459>
4. Gajdzik, B., Wolniak, R., & Grebski, W. W. (2023). Electricity and heat demand in steel industry technological processes in Industry 4.0 conditions. *Energies*. <https://doi.org/10.3390/en16020787>

5. Shi, Y., Yang, F., Zhang, Y., Wang, S., & Han, J. (2024). Empirical analysis of monthly electricity consumption prediction in manufacturing industry using machine learning techniques. *Applied Mathematics and Nonlinear Sciences*. <https://doi.org/10.2478/amns-2024-3142>
6. Chuan, Z. L., et al. (2024). Sustainable energy management: Artificial intelligence-based electricity consumption prediction in limited dataset environment for industry applications. *MATEMATIKA*. <https://doi.org/10.11113/matematika.v40.n3.1594>
7. Zhan, X., et al. (2024). Predicting industrial electricity consumption using industry-geography relationships: A graph-based machine learning approach. *Energies*. <https://doi.org/10.3390/en17174296>
8. Smith, J., Brown, K., & Taylor, L. (2023). Energy consumption forecasting in industrial sectors using machine learning techniques: A comprehensive review. *Energy Reports*, 9(2), 1345–1362.
9. Wang, H., Liu, X., & Zhao, M. (2023). A hybrid model for industrial energy consumption prediction based on LSTM and ARIMA. *Journal of Industrial Energy Management*, 15(3), 876–892.
10. Garcia, P., Torres, L., & Fernandez, A. (2023). Predicting energy consumption in manufacturing industries using ensemble learning techniques. *International Journal of Energy Optimization*, 18(4), 675–690.
11. Patel, S., Kumar, R., & Singh, P. (2023). Sustainable energy management in Industry 4.0: Challenges and opportunities. *Renewable and Sustainable Energy Studies*, 20(2), 345–362.
12. Ansar, T., & Ashraf, W. M. (2025). Comparison of Kolmogorov–Arnold networks and multi-layer perceptron for modelling and optimisation analysis of energy systems. *Energy AI*. <https://doi.org/10.1016/j.egyai.2025.100473>
13. Ashraf, W. M., & Dua, V. (2023). Artificial intelligence-driven smart operation of large industrial complexes supporting the net-zero goal: Coal power plants. *Digital Chemical Engineering*. <https://doi.org/10.1016/j.dche.2023.100119>

14. Mughees, N., Jaffery, M. H., Mughees, A., Ansari, E. A., & Mughees, A. (2023). Reinforcement learning-based composite differential evolution for integrated demand response scheme in industrial microgrids. *Applied Energy*, 342, 121150.
15. François-Lavet, V., Taralla, D., Ernst, D., & Fonteneau, R. (2016). Deep reinforcement learning solutions for energy microgrids management. In *European Workshop on Reinforcement Learning (EWRL 2016)*.
16. Chen, P., Liu, M., Chen, C., & Shang, X. (2019). A battery management strategy in microgrid for personalized customer requirements. *Energy*, 189, 116245.
17. Nakabi, T. A., & Toivanen, P. (2021). Deep reinforcement learning for energy management in a microgrid with flexible demand. *Sustainable Energy, Grids and Networks*, 25, 100413.
18. Ji, Y., Wang, J., Xu, J., Fang, X., & Zhang, H. (2019). Real-time energy management of a microgrid using deep reinforcement learning. *Energies*, 12, 2291.
19. Lee, S., Seon, J., Sun, Y. G., Kim, S. H., Kyeong, C., Kim, D. I., & Kim, J. Y. (2023). Novel architecture of energy management systems based on deep reinforcement learning in microgrid. *IEEE Transactions on Smart Grid*.
20. Ahmed, I., Pedersen, A., & Mihet-Popa, L. (2024). Smart microgrid optimization using deep reinforcement learning by utilizing the energy storage systems. In *Proceedings of the 4th International Conference on Smart Grid and Renewable Energy (SGRE 2024)* (pp. 1–7). IEEE.
21. Fujimoto, S., Hoof, H., & Meger, D. (2018). Addressing function approximation error in actor-critic methods. In *Proceedings of the International Conference on Machine Learning (ICML)* (pp. 1587–1596). PMLR.