



Predictive Data Analytics for Fault Diagnosis and Energy Optimization in Industrial IoT Environments

Dina Fallah¹, Bushra Jabbar Abdul-Kareem², Nada Mohammed Murad³, Ammar Falih Mahdi⁴, Ola Janan^{5*}, Siti Sarah Maidin^{6, 7, 8}

¹Al-Turath University, Baghdad, Iraq

²Al-Mansour University College, Baghdad, Iraq

³Al-Mamoon University College, Baghdad, Iraq

⁴Al-Rafidain University College, Baghdad, Iraq

⁵*Madenat Alelem University College, Baghdad, Iraq

⁶Centre for Data Science and sustainable Technologies, Faculty of Data Science and Information Technology, INTI, International University, Nilai, Negeri Sembilan, Malaysia

⁷Department of IT and Methodology, Wekerle Sandor Uzleti Foiskola, Budapest, Hungary

⁸Faculty of Liberal Arts, Shinawatra University, Thailand

*Corresponding author Email: ola.jinan@mauc.edu.iq

The manuscript was received on 20 November 2024, revised on 29 December 2024, and accepted on 30 March 2025, date of publication 10 May 2025

Abstract

The fusion of predictive maintenance with energy optimization represents a critical advance for intelligent Industrial Internet of Things (IIoT) systems. In response to the growing industrial demand for highly reliable and efficient operations, this study introduces and validates a unified framework that couples fault diagnosis via deep learning with energy management via reinforcement learning. We utilize a Convolutional Neural Network-Long Short-Term Memory (CNN-LSTM) architecture for multivariate fault detection, which demonstrates superior classification accuracy and robustness against data incompleteness. Simultaneously, a Deep Q-Network (DQN) performs dynamic energy scheduling based on predicted system health, achieving substantial energy reductions without compromising task deadlines. Extensive experimental results from real-world industrial datasets and simulations confirm the integrated framework's superiority over conventional approaches in both diagnostic precision and energy efficiency. Key performance indicators, including inference speed and cross-validation, affirm its suitability for real-time industrial applications. This work demonstrates that integrating predictive analytics into intelligent control paradigms is crucial for improving the reliability and sustainability of modern IIoT systems and offers a replicable blueprint for developing next-generation smart manufacturing solutions.

Keywords: Task Offloading Industrial IoT, Predictive Maintenance, Fault Diagnosis, Energy Optimization, Deep Learning.

1. Introduction

The advent of the Industrial Internet of Things (IIoT) has triggered a paradigm shift in industrial operations, enabling intelligent automation, remote supervision, and real-time decision-making in complex environments. As sensor networks and edge computing platforms become ubiquitous, industries now generate vast streams of operational data. However, harnessing this data effectively is constrained by two critical and often interconnected challenges: ensuring operational reliability through proactive fault management and optimizing energy consumption for economic and environmental sustainability. Unplanned system downtime and energy inefficiencies not only disrupt production but also escalate operational costs and accelerate equipment degradation [1], [2]. While a growing body of research has addressed these challenges, solutions have typically evolved in isolation. On one hand, advancements in predictive maintenance have improved fault diagnosis using machine learning (ML) and deep learning (DL) models [4], [5]. On the other hand, separate research has focused on energy optimization, using methods to minimize waste and support intelligent maintenance strategies [3]. Despite these parallel advances, a significant gap persists: the lack of integrated frameworks that fuse predictive fault analytics with



Copyright © Authors. This is an open access article distributed under the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

real-time energy control. Existing solutions are often implemented in siloed subsystems, neglecting the crucial feedback loop between a system's health status and its energy consumption profile.

This fragmentation limits the effectiveness of current technologies. Conventional rule-based and statistical monitoring systems struggle to adapt to the dynamic nature of industrial settings, which involve sensor drift, component aging, and variable workloads. While sophisticated temporal models like CNNs and LSTMs have shown promise in classifying faults from time-series data, they rarely incorporate energy-saving optimization layers [6]. Consequently, a system might accurately predict an impending failure but continue to operate in an energy-inefficient mode, or a predictive maintenance alert might trigger actions without considering the most energy-opportune schedule. This disconnect represents a missed opportunity for holistic optimization. This paper addresses these shortcomings by proposing a unified predictive analytics framework that simultaneously performs fault diagnosis and energy optimization in IIoT environments. The novelty of our approach lies in its co-dependent architecture, where a hybrid Convolutional Neural Network-Long Short-Term Memory (CNN-LSTM) model provides robust fault predictions that directly inform an intelligent energy management module built on Deep Q-Network (DQN) reinforcement learning. These components interact dynamically, enabling proactive adjustments to operational parameters based on predicted fault probabilities while ensuring energy-efficient task execution. The main contributions of this work involve the development of an integrated DL-based framework for accurately diagnosing multivariate faults in real-time IIoT networks; the design of an adaptive, energy-aware control policy using deep reinforcement learning that responds dynamically to system health predictions; and a comprehensive validation of the coupled framework's efficacy on industrial benchmark datasets, which demonstrates significant improvements in both diagnostic accuracy and energy conservation over conventional, non-integrated models. By bridging the gap between diagnostics and optimization, this research seeks to enhance the dependability, resilience, and sustainability of IIoT-enabled smart manufacturing systems.

2. Literature Review

The exponential growth of Industrial IoT (IIoT) systems has spurred a significant body of research into data-driven methods for enhancing operational intelligence. Predictive maintenance and energy management have emerged as two cornerstone applications driving the smart manufacturing era. However, a comprehensive review of the current literature reveals that while substantial progress has been made in each domain, the critical link between predictive fault diagnosis and energy-aware operational control remains underdeveloped and fragmented.

2.1. Advances in Predictive Fault Diagnosis

Modern fault diagnosis has increasingly shifted from reactive to predictive strategies, leveraging the power of machine learning. Early approaches often relied on statistical methods, but these have proven insufficient for the dynamic and complex nature of IIoT environments. More recently, sophisticated models have demonstrated significant promise. Graph-based modeling, for instance, has been shown to be an effective tool for diagnosing faults in large, networked systems like energy grids. Zhang et al. [9] proposed a graph learning method to model topological relationships in energy networks, improving failure localization. To handle heterogeneous data streams from wireless sensor networks (WSNs), researchers have developed advanced classification models. Lavanya et al. [10] introduced a tuned classification system that enhanced diagnostic accuracy in IoT-enabled WSNs. Similarly, ensemble methods that fuse multiple base classifiers have been used to improve the robustness of fault detection in building energy systems [11]. However, a common limitation of these high-accuracy models is their computational intensity, which often makes them unsuitable for resource-constrained edge devices where energy conservation is paramount. Furthermore, many of these models struggle with real-world data challenges, such as the class imbalance and data incompleteness highlighted by Lin and Jamrus [15].

2.2. Data-Driven Energy Management and Optimization

Parallel to advances in diagnostics, significant research has focused on optimizing energy consumption in industrial systems. Big data platforms have been developed to create scalable ecosystems for predictive maintenance, but these systems typically prioritize data volume and processing over real-time energy management [12]. Deep learning models have also been applied to improve power generation forecasting by identifying faulty records in IoT data streams, leading to more reliable energy output predictions [13]. Other studies have focused on fault prediction for specific assets, such as smart grid equipment, achieving strong classification performance [14]. However, these diagnostic insights are rarely used to inform dynamic energy control strategies. The models can predict a fault but lack the capability to translate that prediction into an adaptive operational plan that minimizes energy consumption under uncertain conditions. While the strategic importance of AI-supported ESG principles in power systems is recognized, the focus has largely remained at a high-level policy level rather than on implementation-ready models for real-time, energy-intelligent diagnostics [18].

2.3. The Research Gap: The Disconnect Between Diagnostics and Optimization

Despite individual advances, a persistent gap exists in the literature: the absence of unified, lightweight frameworks that simultaneously achieve accurate diagnostics and real-time energy optimization. The two fields have evolved largely in parallel. Fault diagnosis systems often overlook the energy implications of their predictions, while energy optimization strategies are typically not informed by the real-time health status of the system. This dichotomy between diagnostic insight and operational action represents a major void in existing paradigms. Several challenges contribute to this gap. As reviewed by Zhao et al. [16], AI-based fault diagnosis systems often suffer from overfitting, poor generalization, and high computational overhead, making deployment in constrained IIoT environments difficult. Early frameworks that attempted to link fault prediction to action relied on simple rule-based automation and gave minimal attention to optimizing energy during corrective measures [17]. More recent anomaly detection tools have improved operational awareness, but they tend to operate reactively rather than predictively [27]. Other complex deep learning methods, while precise, come at the cost of high complexity that is impractical for real-time control [30]. In summary, the literature shows a clear need for an integrated approach. Existing models are often too complex for edge deployment, are not robust against real-world data imperfections, or treat fault diagnosis and energy management as separate tasks. This research aims to fill this critical gap by proposing and validating a framework that explicitly couples predictive diagnostics with adaptive energy control, creating a synergistic system that is both reliable and efficient.

3. Methods

This study employs a hybrid, data-driven methodology to develop and validate an integrated framework for predictive fault diagnosis and energy optimization in Industrial IoT (IIoT) environments. The approach combines deep learning for time-series analysis with reinforcement learning for dynamic control, evaluated through rigorous experimentation on real-world and benchmark industrial datasets. The research is structured into five sequential stages: (1) data acquisition and experimental design, (2) data preprocessing and feature engineering, (3) development of the fault diagnosis module, (4) development of the energy optimization module, and (5) system integration and validation.

3.1. Data Collection and Experimental Design

To ensure the robustness and generalizability of our findings, a comprehensive dataset was curated from multiple sources to reflect diverse operational conditions. The primary data was sourced from three industrial renewable energy plants (wind and solar hybrid) in Northern Europe and was supplemented by public benchmark datasets [8], [13], [30]. This combined dataset comprises over 3.2 million timestamped sensor data entries from 62 IoT nodes covering key physical parameters; qualitative insights from 11 structured interviews with on-site system engineers and energy managers to understand operational context; and 19 detailed fault incident reports collected over a 15-month period to provide ground-truth labels for supervised learning. The experimental setup was subsequently designed to emulate realistic industrial conditions, systematically introducing scenarios involving artificially injected noise, missing data, power fluctuations, and class imbalances to test the framework's resilience and performance under non-ideal circumstances [8], [22].

3.2. Data Preprocessing and Feature Engineering

Raw sensor data is often noisy, inconsistent, and contains missing values. A multi-step preprocessing pipeline was implemented to prepare the data for model training. First, all sensor readings were synchronized using their timestamp keys to create a unified temporal view. An Exponential Moving Average (EMA) filter was then applied to mitigate high-frequency noise. Outliers were identified and handled using the Interquartile Range (IQR) method, and missing values (constituting approximately 4.3% of the dataset) were imputed using a graph-based interpolation technique to preserve spatio-temporal relationships [9]. Following preprocessing, a rich feature set was engineered to capture the complex dynamics of the system. A total of 64 features were extracted, falling into three categories: statistical features (e.g., mean, variance) to summarize data distribution, temporal features (e.g., lag features, rolling windows) to capture time-based dependencies, and spectral features derived from Fourier transforms to analyze frequency-domain characteristics. To reduce dimensionality and prevent model overfitting, Principal Component Analysis (PCA) was then applied to this feature set. The analysis revealed that the first 11 principal components successfully captured 92.7% of the total variance, and these were selected as the final input features for the diagnostic model [15].

3.3. Fault Diagnosis Model

For fault diagnosis, this study adopted a hybrid deep learning architecture that combines a 1D Convolutional Neural Network (CNN) with a Long Short-Term Memory (LSTM) network. This dual approach was chosen for its ability to effectively process complex multivariate time-series data. The CNN layers excel at automatically extracting salient, localized features from raw sensor signals, while the LSTM layers are adept at modeling the temporal dependencies and long-range patterns within these features, which is crucial for predicting system failures over time. Let $X_t \in \mathbb{R}^{n \times m}$ represent the input matrix of sensor readings over time:

$$X_t = [x_{t-n}, x_{t-n+1}, \dots, x_t] \quad (1)$$

The diagnostic model output \hat{y}_t is obtained as:

$$\hat{y}_t = \text{Softmax}(W_l \cdot \text{LSTM}(CNN(X_t)) + b_l) \quad (2)$$

where W_l and b_l are trainable weights and biases. The model achieved F1-score = 0.94, with diagnostic latency ~ 210 ms, significantly outperforming tuned classification models in prior WSN applications [10]. Model ensemble was also tested using AdaBoost and XGBoost for decision refinement [11].

3.4. Energy Optimization via Reinforcement Learning

To optimize energy consumption dynamically, we implemented a Deep Q-Network (DQN), a model-free reinforcement learning algorithm. RL is well-suited for this problem as it allows an agent to learn an optimal control policy through direct interaction with the environment, without needing an explicit system model. The DQN uses a neural network to approximate the optimal action-value function (Q-function), making it capable of handling high-dimensional state spaces.

The problem was formulated as a Markov Decision Process (MDP) with the following components: a) The state vector provides a comprehensive snapshot of the system at time t . It includes the current sensor data, the predicted fault probability vector from the CNN-LSTM module, and the current operational load demand, b) The action space consists of a discrete set of energy reallocation strategies, such as throttling power to non-critical subsystems, deferring high-load tasks, or rerouting processes to healthy components, c) The reward function was engineered to guide the agent toward a multi-objective goal of balancing energy savings, system reliability, and operational performance. It is defined as:

$$R_t = -\alpha \cdot E_t - \beta \cdot (1 - F_t) - \gamma \cdot \mathbb{I}_{\text{Overload}} \quad (3)$$

Where E_t energy consumed at time t , F_t fault probability, $\mathbb{I}_{\text{Overload}}$ overload indicator function.

Tuned hyperparameters: $\alpha = 0.7$, $\beta = 0.2$, $\gamma = 0.1$, selected via grid search over 400 episodes. The optimization led to average energy savings of 17.4%, with less than 2.5% task deadline violation, outperforming traditional rule-based systems [2], [14].

3.5. Integration and Real-Time Execution

The proposed framework is designed as a real-time processing pipeline, as illustrated in Figure 1. The architecture begins with sensor-level data acquisition and flows through a sequence of preprocessing, diagnostic inference, and dynamic energy control. To facilitate modularity and scalability, the diagnostic and optimization subsystems were deployed within a Dockerized microservice architecture. Data is streamed from IoT devices using the lightweight MQTT protocol and stored in a time-series database (InfluxDB) optimized for high-throughput industrial data.

The CNN-LSTM module serves as the primary diagnostic engine, analyzing multivariate sensor inputs to detect anomalies and anticipate fault conditions. These predictive outputs feed directly into the DQN scheduler, which then selects the most energy-efficient control action based on the predicted system health state. This creates a dynamic feedback loop, ensuring that both fault predictions and energy decisions adapt to real-time operational changes. For operational oversight, a visualization and alerting dashboard was integrated using Grafana, enabling facility operators to monitor system status and intervene when diagnostic confidence thresholds exceed 0.85. To enhance model interpretability, particularly in complex failure modes, Spatial-Temporal Graph Neural Networks (GNNs) were employed for anomaly explanation in cases of multivariate drift or when missing data surpassed 8% [7], aligning with recent methods for robust inference in uncertain environments [9].




Fig 1. Architecture of the integrated predictive analytics framework for fault diagnosis and energy optimization in IIoT systems

3.6. Research Hypotheses and Validation

This study was designed to test two primary hypotheses:

H1: An integrated framework that uses deep learning-based fault predictions to inform a reinforcement learning-based energy optimization policy can significantly reduce overall energy consumption without sacrificing diagnostic accuracy or operational performance.

H2: A dual-model architecture (CNN-LSTM and DQN) is more robust and resilient to real-world operational challenges, such as noise and missing data, than single-purpose diagnostic or optimization pipelines.

To rigorously validate these hypotheses, a multi-faceted evaluation strategy was employed. The generalization performance and stability of the CNN-LSTM model were assessed using a 5-fold cross-validation approach. To test the framework's robustness under uncertainty, Monte Carlo simulations ($n=100$) were conducted with varying levels of injected noise and missing data. The statistical significance of performance differences between our proposed framework and baseline models was assessed at a confidence level of $p < 0.05$ across all major performance indicators [12], [16].

4. Result and Discussion

4.1. Diagnostic Model Performance Evaluation

The diagnosis metrics (accuracy, speed, and classification) of machine learning models used to detect failures in problems of the industrial Internet of things (IIoT) based on the Industry 4.0. The models involve Long Short-Term Memory (LSTM), hybrid CNN-LSTM, XGBoost, AdaBoost, and specially tailored Ensemble Model. The algorithms were trained and evaluated based on a large multivariate time-series dataset from real industrial sensor installations. Performance was compared using accuracy, F1-score, precision, recall, and latency, where latency is measured in milliseconds. The objective is to investigate whether the deep learning structures, or the sequential hybrid networks, are more stable and more accurate compared with traditional or single models. It is important for IIoT applications to have such ability; immediate and accurate fault detection is also a key-point to minimize downtime and to increase energy efficiency.




Fig 2. Performance metrics of fault diagnosis models

The results in Figure 2 indicate that the Ensemble Model or combination of the three individual models was superior over all measures where it reached 96.5% in accuracy and an F1-score of 0.97. But its latency at 260 milliseconds was the longest of all models, and might not be as well-suited for time-sensitive uses. Given the trade-offs between accuracy and speed, the CNN-LSTM was the best performing model, and it achieved an accuracy of 95.8% with a latency of 230 ms. Classical boosting techniques such as XGBoost (F1-score 0.89) and AdaBoost (F1-score 0.88) performed worse respectively. (XGBoost was the fastest, though by only 140 ms, but it had less recall of 0.88, which is particularly important for avoiding false negatives when fault detection might miss something.) These results justify the use of CNN-LSTM as main engine for diagnosis, between latency and accuracy trade-off, when a fault recognition generalization across heterogeneous IIoT environments is required.

4.2. Energy Optimization Strategy Results

The efficiency of energy optimization techniques used with predictive fault detection models is analyzed. It was compared to heuristic optimization, a genetic algorithm, regular Q-learning and to a rule-based baseline using a Deep Q Network (DQN). Performance was quantified with the average energy percentage spared, percentage of missed jobs deadlines, total number of learning episodes, and average reward scores reached throughout the reinforcement learning process. Evaluation was performed in a virtual-real industrial testbed with dynamic task scheduling, varying load conditions and varying system health states, based on predictive diagnostics. We aim to understand how much energy optimization techniques can reduce energy consumption while respecting operational constraints (and specifically real-time deadlines) that are crucial for system performance and reliability preservation.




Fig 3. Evaluation of energy optimization methods in IIoT environment

The DQN as shown in Fig. 3 achieves better optimization performance with a lower energy consumption of 17.4% at a low task deadline miss rate of 2.4%. This further confirms the model's capability to trade-off energy efficiency and real-time system responsiveness. The heuristic and baseline methods, however, obtained much less energy savings (10.4% and 8.3%, respectively) and higher miss rates (5.6% and 6.8%). The standard Q-learning was also tested, producing a comparable 14.2% energy savings with a 3.7% miss rate, but lagged behind the DQN in terms of total reward earned. The genetic algorithm performed in middle, stable but not adaptive in dynamic environment. The high value of reward score for DQN makes it more effective in learning in different scenarios. DQN, in general, offers a promising approach to energy-aware task scheduling in IIoT-based systems complemented with predictive diagnostics.

4.3. Cross-Validation of CNN-LSTM Diagnostic Model

The CNN-LSTM architecture was cross-validated with a 5-fold partitioning to guarantee model robustness. The folds were drawn randomly over temporal segments to evaluate generalization across temporal windows. Performance measures are classification accuracy, F1-score, and the loss functions. This is important for industrial applications where future fault patterns are likely to be different to the present under external or seasonal conditions. CV is useful to assess how much the trained model is reliable and consistent between different operating conditions.




Fig 4. Five-fold cross-validation results for CNN-LSTM model

Highly stable performance was also observed for the 5-fold cross validation evaluations where the CNN-LSTM was able to score very high accuracies, ranging from 95.1 to 95.8% (Figure 4). The F1-score also exhibited a narrow range (0.95–0.96), which indicates that the model was not overfitting or unstable. Losses consistently decreased from 0.212 in Fold 1 to 0.200 in Fold 5, which indicated that the convergence was enhanced during training and validation (the lower the better). These findings verify that the architecture not only is accurate but also stable across different data segments, and thus is amenable for long-term deployment in IIoT conditions characterized by temporal variation and non-stationarity. The marginal difference in loss also indicates that the architecture is robust against change in fault distribution over time.

4.4. Impact of Missing Data on Diagnostic Accuracy

For noise and unreliable system simulation in practice, the CNN-LSTM model was tested with different ratios of the missing data, which were artificially injected. This is particularly crucial for IIoT scenarios of sensors failing or experiencing communication latency. Accuracy was estimated at different rates of missing observations between 0% (reference) and 20%. In this test, the fault tolerance of the diagnostic system was tested, too, and it was evaluated, how much does the diagnostic system keep its performance if the information is not complete.




Fig 5. Diagnostic accuracy of CNN-LSTM under missing data conditions

The model exhibited a graceful deteriorative performance as the ratio of missing data increased. With 5% missing input, starting from a baseline of 95.8% accuracy (with full input), the model still performed at 94.5% accuracy. The accuracy decreased slightly with 10% and 15% missing data, to 93.1% and 91.2%, respectively. The performance was < 90% only at the largest tested threshold of 20%, and thus high resilience. The findings confirm the suitability of the CNN-LSTM in realistic scenarios with sensor malfunctioning and data

degradation. Its robustness can also be ascribed to its sequential memory structure as well as the precursory prepossessing of data in the method.

4.5. Real-Time Inference Speed and Resource Efficiency

Predictive models should be very accurate subject to strong constraints on time and computational complexity in IIoT environments. Inference throughput, latency, CPU usage, and memory consumption were evaluated for a variety of core designs including LSTM, CNN-LSTM, XGBoost, Q-Learning, and DQN. These measurements were obtained with the deployment on industrial grade edge devices featuring quad-core ARM CPUs and 8GB of RAM. It is important to gauge the inference performance, as this will be important for verifying the suitability of a model in edge-based applications, where there are scarce hardware resources, and energy efficiency is crucial. A good model should at the same time provide a high throughput, with low latency and low resource consumption, such that the reliable, scalable and real time operation can be achieved in the FD (fault detection) and EO (energy optimization) workflow.




Fig 6. Real-time inference speed and resource utilization across architectures

Also, the best performing model XGBoost had the best inference throughput (102 samples/s), and the lowest average latency (125 ms), suggesting that it is a cost-effective alternative for fast inference when resource constraints are tight. However, lower diagnostic accuracy of EIOM restricts its application in the prediction of critical faults. Although the CNN-LSTM outperformed other approaches in accuracy, it also recorded the largest memory utilization, 2.2 GB and second-highest CPU utilization to imply that a local level optimization was necessary. The DQN was also efficient in terms of resource utilization: it utilized 2.6 GB of RAM and 68% of the CPU, which is reasonable given that it is a dual learning method across states and rewards. The results also show that LSTM achieved a good trade-off between speed (latency and memory) and performance, so we consider it as a beneficial fallback model. These findings highlight the near-equivalence of computational complexity and diagnostic fidelity in real-time IIoT use.

4.6. System-Wide Energy Consumption: Before vs After Optimization

In order to estimate the energy saving of the optimization model, we compared PARCEL energy consumption with that of a baseline on 5 main subsystems: power conversion, sensor bus, data aggregation, edge processing, and cloud synchronization. Data was obtained with power monitoring modules placed in the operating control units over a 30-day assessment window. This evaluation compared the energy signatures with and without integrating the novel DQN-based optimization module. The aim was to investigate the effects of intelligent energy scheduling on the performance of individual subsystems and the total power consumed, while maintaining the reliability of service and real-time monitoring and analysis metrics.




Fig 7. System-wide energy usage before and after applying optimization

Post-optimization conclusions reveal large energy reductions in all observed subsystems. Significant reductions were seen in Power Conversion and Edge Components of 2.3 kWh and 2.2 kWh. This is consistent with the focused scheduling of high-load activities during anticipated low-fault windows, allowed by the DQN architecture. The sensor bus and data aggregations devices, have already

being optimized for low energy consumption, but still made over 0.6 kWh and 1.4 kWh, respectively. There was a significant 2.2 kWh reduction of cloud synchronization, due to an improved data offloading strategy. Overall, these results verify that the incorporation of fault-aware scheduling into energy operations not only decreases the overall power consumption but also leads to a balanced and efficient distribution of the energy across the entire IIoT environment.

4.7. Discussion

The experimental results of this study provide strong evidence for the efficacy of integrating predictive fault diagnostics with reinforcement learning-based energy optimization in Industrial IoT (IIoT) environments. The findings confirm our primary hypotheses: that a unified framework can significantly enhance system reliability and energy efficiency simultaneously. The high performance of the CNN-LSTM diagnostic model, coupled with the effectiveness of the DQN energy scheduler, demonstrates that creating an intelligent feedback loop between a system's predicted health and its operational strategy is both technologically feasible and practically beneficial.

4.7.1. Interpretation of Principal Findings

The core success of this work lies in its validation of a synergistic, dual-model architecture that successfully marries predictive insight with intelligent action. The CNN-LSTM model's ability to achieve 95.8% accuracy substantiates its suitability for navigating the complexities of multivariate industrial time-series data. This high level of diagnostic precision is the foundation upon which the entire framework's value is built, as it provides the reliable, real-time health assessments necessary for confident decision-making. The model's robust performance serves as a testament to the power of hybrid deep learning architectures in extracting both spatial and temporal features from noisy sensor streams. More importantly, the framework translates this diagnostic accuracy into tangible operational gains. The DQN scheduler's capacity to leverage these predictions to achieve a 17.4% reduction in energy consumption while maintaining a task deadline miss rate below 2.5% directly supports our central hypothesis. This outcome illustrates a strategic shift from a reactive or siloed operational paradigm to a proactive and holistic one, where resource management is intelligently and dynamically informed by predictive analytics. Furthermore, the framework's resilience, maintaining over 91% accuracy with 15% missing data, further underscores its potential for reliable deployment in real-world industrial settings where sensor failures and communication gaps are common occurrences.

4.7.2. Comparison with Prior Work

When contextualized within the existing literature, the contribution of our integrated approach becomes clear, particularly in its departure from single-focus solutions. Previous studies in fault diagnosis, such as those by Lavanya et al. [10] and Han et al. [11], have made strides in enhancing classification accuracy but have largely overlooked the critical trade-offs between performance, computational cost, and energy awareness. Our framework explicitly addresses this by not only achieving high diagnostic accuracy but also ensuring the model is computationally viable for edge deployments, where resources are often constrained. Similarly, our work advances beyond prior research in energy management, which has often lacked a direct, real-time link to system health. Big data ecosystems like the one proposed by Yu et al. [12] have concentrated on scalable infrastructure for fault detection without closing the loop to enable responsive energy optimization. Our DQN agent's performance surpasses that of traditional heuristic and evolutionary methods, validating its effectiveness as a truly adaptive scheduler. It directly responds to diagnostic information, addressing a capability gap left by prior diagnostic algorithms for smart grids [14] and creating an actionable bridge between knowing a fault may occur and knowing how to operate more efficiently because of that knowledge.

4.7.3. Limitations and Future Research Directions

Despite the promising results, several limitations must be acknowledged to frame the scope of this work accurately. First, while the framework was rigorously tested on real-world datasets, the experimental setting was ultimately a controlled simulation. Future work must focus on deploying and validating the system in a live, heterogeneous industrial environment to assess its performance against unforeseen operational challenges, complex multi-system interactions, and policy-driven constraints that cannot be fully replicated. Second, the DQN agent, while effective, requires a significant number of training episodes for convergence. This training overhead could be a practical barrier to rapid deployment or adaptation in dynamic environments. Future iterations could explore transfer learning or federated reinforcement learning to accelerate this process, allowing models to learn from decentralized data sources without compromising privacy [13]. Furthermore, the model itself presents avenues for algorithmic enhancement. Our approach could be improved to better handle extremely rare fault types, as minor performance fluctuations were observed for highly imbalanced classes a known challenge in industrial datasets [15]. This limitation could be mitigated by incorporating generative adversarial networks (GANs) to produce synthetic data for minority classes, thereby creating a more balanced training environment. Finally, while our framework represents a methodological leap from reactive tools [27] to predictive control, its "black box" nature could hinder adoption. Integrating explainable AI (XAI) techniques to provide transparent justifications for both diagnostic and control decisions would likely increase trust and empower system operators to work collaboratively with the AI.

4.7.4. Broader Implications

The findings of this study have significant implications for the future of smart manufacturing and the push toward sustainable industrial automation. By demonstrating tangible energy savings across critical IIoT subsystems, this work provides a technical blueprint that directly aligns with and supports broader ESG (Environmental, Social, and Governance) principles in the energy and manufacturing sectors [18]. It moves beyond high-level policy discussions to offer a practical, implementation-ready methodology for creating industrial processes that are not only more productive and reliable but also more environmentally and economically sustainable. Ultimately, this research lays the groundwork for more advanced, fully autonomous systems. The successful integration of intelligent diagnostics with adaptive control provides a scalable and repeatable methodology for enhancing the resilience and efficiency of next-generation industrial infrastructure. It opens up future research avenues into more complex systems that incorporate real-time online learning, dynamic edge-cloud federated intelligence, and adaptive meta-controllers. Such systems could optimize policies based not only on system health but also on external dynamic factors like fluctuating energy prices, supply chain disruptions, or evolving sustainability targets, paving the way for truly intelligent and resilient industrial automation.

5. Conclusion

This study successfully developed and validated a unified framework that integrates predictive fault diagnostics with intelligent energy optimization for Industrial Internet of Things (IIoT) systems. By coupling a CNN-LSTM diagnostic model with a DQN energy scheduler, we have demonstrated that it is not only technologically feasible but also practically significant to create a synergistic feedback loop between a system's predicted health and its operational strategy. The results confirm our primary hypothesis: this integrated approach remarkably enhances both system reliability and energy efficiency, outperforming traditional methods that address these challenges in isolation. The core contribution of this work is a major step forward in IIoT system design, enabling a strategic shift from reactive responses to proactive, predictive control. The framework provides a blueprint for aligning real-time system performance with key business drivers such as cost reduction, operational uptime, and sustainability. The proven robustness of the CNN-LSTM model against noisy and incomplete data, combined with the DQN scheduler's ability to adapt energy consumption without compromising performance, establishes a new benchmark for intelligent industrial automation.

While the framework performed exceptionally well, future work should focus on validating its scalability and generalization in diverse, real-world industrial deployments. Key research directions include the application of federated learning to reduce training overhead in distributed environments and the integration of explainable AI (XAI) to enhance operator trust and adoption. Further improvements could also be made in handling rare fault types through generative modeling techniques. In conclusion, the capacity to predict and prevent failures while simultaneously optimizing energy in real time is a foundational requirement for the next generation of data-driven, autonomous industrial systems. This research provides a reproducible and scalable methodology to meet this demand, laying the groundwork for future investigations into multi-objective optimization, adaptive control, and the development of truly resilient and sustainable smart factories.

References

- [1] Sharma, N., et al., Energy-Efficient and QoS-Aware Data Routing in Node Fault Prediction Based IoT Networks. *IEEE Transactions on Network and Service Management*, 2023. 20(4): p. 4585-4599.
- [2] Atassi, R.A., F., Predictive Maintenance in IoT: Early Fault Detection and Failure Prediction in Industrial Equipment. *Journal of Intelligent Systems and Internet of Things*, 2023. 9(2): p. 231-238.
- [3] Zhang, J., et al., Fault diagnosis and intelligent maintenance of industry 4.0 power system based on internet of things technology and thermal energy optimization. *Thermal Science and Engineering Progress*, 2024. 55: p. 102902.
- [4] Li, X., et al., Energy-Propagation Graph Neural Networks for Enhanced Out-of-Distribution Fault Analysis in Intelligent Construction Machinery Systems. *IEEE Internet of Things Journal*, 2025. 12(1): p. 531-543.
- [5] Nathiya, N., C. Rajan, and K. Geetha, A hybrid optimization and machine learning based energy-efficient clustering algorithm with self-diagnosis data fault detection and prediction for WSN-IoT application. *Peer-to-Peer Networking and Applications*, 2025. 18(2): p. 13.
- [6] Yao, Y., et al., Small-Batch-Size Convolutional Neural Network Based Fault Diagnosis System for Nuclear Energy Production Safety With Big-Data Environment. *International Journal of Energy Research*, 2020. 44(7): p. 5841-5855.
- [7] Zhang, J., Y. Cheng, and X. He, Fault Diagnosis of Energy Networks Based on Improved Spatial-Temporal Graph Neural Network With Massive Missing Data. *IEEE Transactions on Automation Science and Engineering*, 2024. 21(3): p. 3576-3587.
- [8] Jovicic, E., et al., Publicly Available Datasets for Predictive Maintenance in the Energy Sector: A Review. *IEEE Access*, 2023. 11: p. 73505-73520.
- [9] Zhang, J., Y. Cheng, and X. He, Fault Diagnosis of Energy Networks: A Graph Embedding Learning Approach. *IEEE Transactions on Instrumentation and Measurement*, 2022. 71: p. 1-11.
- [10] Lavanya, S., et al., A Tuned classification approach for efficient heterogeneous fault diagnosis in IoT-enabled WSN applications. *Measurement*, 2021. 183: p. 109771.
- [11] Han, H., et al., Ensemble learning with member optimization for fault diagnosis of a building energy system. *Energy and Buildings*, 2020. 226: p. 110351.
- [12] Yu, W., et al., A Global Manufacturing Big Data Ecosystem for Fault Detection in Predictive Maintenance. *IEEE Transactions on Industrial Informatics*, 2020. 16(1): p. 183-192.
- [13] Rani, N.C., et al. Power Generation Forecasting Through IoT-Driven Fault Detection Using Deep Learning. in 2024 International Conference on Advancement in Renewable Energy and Intelligent Systems (AREIS). 2024.
- [14] Ming Gao, F.Z., Research on Fault Diagnosis and Prediction Algorithms for Power Equipment in Smart Grids. *Journal of Electronics and Information Science*, 2024. 9(2): p. 114-119.
- [15] Lin, K.-Y. and T. Jamrus, Industrial data-driven modeling for imbalanced fault diagnosis. *Industrial Management & Data Systems*, 2024. 124(11): p. 3108-3137.
- [16] Zhao, Y., et al., Artificial intelligence-based fault detection and diagnosis methods for building energy systems: Advantages, challenges and the future. *Renewable and Sustainable Energy Reviews*, 2019. 109: p. 85-101.
- [17] Wang, C., H.T. Vo, and P. Ni. An IoT Application for Fault Diagnosis and Prediction. in 2015 IEEE International Conference on Data Science and Data Intensive Systems. 2015.
- [18] Li, Q., et al., ESG guidance and artificial intelligence support for power systems analytics in the energy industry. *Scientific Reports*, 2024. 14(1): p. 11347.
- [19] J. Prayitno, B. Saputra, and A. Kumar, "Emotion Detection in Railway Complaints Using Deep Learning and Transformer Models: A Data Mining Approach to Analyzing Public Sentiment on Twitter," *Journal of Digital Society*, vol. 1, no. 2, pp. 1-14, 2025, doi.org/10.63913/jds.v1i2.6.
- [20] A. D. Buchdadi and A. S. M. Al-Rawahna, "Anomaly Detection in Open Metaverse Blockchain Transactions Using Isolation Forest and Autoencoder Neural Networks," *International Journal Research on Metaverse*, vol. 2, no. 1, pp. 24-51, 2025, doi: 10.47738/ijrm.v2i1.20.

[21] Y. Durachman and A. W. Bin Abdul Rahman, "Clustering Student Behavioral Patterns: A Data Mining Approach Using K-Means for Analyzing Study Hours, Attendance, and Tutoring Sessions in Educational Achievement," *Artificial Intelligence in Learning*, vol. 1, no. 1, pp. 35–53, 2025, doi: 10.63913/ail.v1i1.5.

[22] Wang, H., et al., Fault Diagnosis Algorithm Based on Power Outage Data in Power Grid. EAI Endorsed Transactions on Energy Web, 2023. 10.

[23] Si, J., Y. Li, and S. Ma, Intelligent Fault Diagnosis for Industrial Big Data. *Journal of Signal Processing Systems*, 2018. 90(8): p. 1221-1233.

[24] A. B. Prasetio, M. Aboobaider, and A. Ahmad, "Assessing Geographic Disparities in Campus Killings: A Data Mining Approach Using Cluster Analysis to Identify Demographic Patterns and Legal Implications," *Journal of Cyber Law*, vol. 1, no. 1, pp. 1–21, 2025, doi.org/10.63913/jcl.v1i1.1.

[25] S. Y. Baroud, N. A. Yahaya, and A. M. Elzamly, "Cutting-Edge AI Approaches with MAS for PdM in Industry 4.0: Challenges and Future Directions," *Journal of Applied Data Sciences*, vol. 5, no. 2, pp. 455–473, 2024, doi: 10.47738/jads.v5i2.196.

[26] E. D. Lusiana, S. Astutik, Nurjannah, and A. B. Sambah, "Using Machine Learning Approach to Cluster Marine Environmental Features of Lesser Sunda Island," *Journal of Applied Data Sciences*, vol. 6, no. 1, pp. 247–258, 2025, doi: 10.47738/jads.v6i1.478.

[27] Piscitelli, M.S., et al., A data analytics-based tool for the detection and diagnosis of anomalous daily energy patterns in buildings. *Building Simulation*, 2021. 14(1): p. 131-147.

[28] A. Wang and Z. Qin, "Development of an IoT-Based Parking Space Management System Design," *International Journal for Applied Information Management*, vol. 3, no. 2, pp. 91–100, 2023, doi: 10.47738/ijaim.v3i2.54.

[29] R. Nagarajan, M. Batumalay, and Z. Xu, "IoT based Intrusion Detection for Edge Devices using Augmented System," *Journal of Applied Data Sciences*, vol. 5, no. 3, pp. 1412–1423, 2024, doi: 10.47738/jads.v5i3.358.

[30] Liu, Y., et al. Power Equipment Fault Diagnosis Method Based on Energy Spectrogram and Deep Learning. *Sensors*, 2022. 22, DOI: 10.3390/s22197330.

[31] Hajji, M., et al., Reducing neural network complexity via optimization algorithms for fault diagnosis in renewable energy systems. *Ain Shams Engineering Journal*, 2024. 15(12): p. 103086.

[32] B. H. Hayadi and I. M. M. El Emary, "Enhancing Security and Efficiency in Decentralized Smart Applications through Blockchain Machine Learning Integration," *Journal of Current Research in Blockchain*, vol. 1, no. 2, pp. 139–154, 2024, doi: 10.47738/jcrb.v1i2.16.

[33] A. R. Hananto and B. Srinivasan, "Comparative Analysis of Ensemble Learning Techniques for Purchase Prediction in Digital Promotion through Social Network Advertising," *Journal of Digital Market and Digital Currency*, vol. 1, no. 2, pp. 125–143, 2024, doi: 10.47738/jdmdc.v1i2.7.

[34] K. Y. Tippayawong, "Construction of Enterprise Logistics Decision Model Based on Supply Chain Management," *International Journal of Informatics and Information Systems*, vol. 6, no. 4, pp. 181–188, 2023, doi: 10.47738/ijiis.v6i4.179.

[35] Buhari Dogan, N., Nketiah, E., Ghosh, S., & Nassani, A. A. (2025). The impact of the green technology on the renewable energy innovation: Fresh pieces of evidence under the role of research & development and digital economy. *Renewable and Sustainable Energy Reviews*, 210, 115193. <https://doi.org/10.1016/j.rser.2024.115193>