Minimax Theory and its Applications

Volume 11 (2026), No. 2s, 461 to 475

AI-Driven Approaches to Enhancing Energy Efficiencyin IoT Networks for Sustainable Performance

Delshi Howsalya Devi R[1], Gunapriya D[2], Sivaranjani S[3], Suganya N[4], Murugesan M[5*], Anand K[6]

[1]Department of Artificial Intelligence and Data Science, Karpaga Vinayaga College ofEngineering and Technology, Tamil Nadu, India. Email: delshi@rocketmail.com
[2]Department of Electrical and Electronics Engineering, Sri Eshwar College of Engineering,Coimbatore, India Email: gunapriya.d@sece.ac.in
[3]Department of Electrical and Electronics Engineering, Sri Krishna College of Engineeringand Technology, Coimbatore, India. Email: sivaranjanis@skcet.ac.in
[4]Department of Computer Science and Business Systems, Dr. N.G.P. Institute ofTechnology, Coimbatore, India Email:nrsuganya@gmail.com
[5]Department of Electrical and Electronics Engineering, Karpagam Institute of Technology,Coimbatore, Tamil Nadu, India. Email: murugesan.kec@gmail.com
[6]Department of Electronics and Communication Engineering, Chennai Institute ofTechnology, Chennai, India Email: anandped2012@gmail.com

Abstract
Energy efficiency in Internet of Things (IoT) networks represents a criticalchallenge as deployments scale exponentially, with battery-powered devicesrequiring sophisticated optimization strategies to minimise operational costsand environmental impact. This study develops and validates artificial intelligence(AI)-driven methods for energy optimization across heterogeneousIoT environments. We collected 2.4 million temporal data points from 150devices spanning smart homes, industrial facilities, and healthcare systemsover a six-month observation period (January–June 2024). Three algorithmicparadigms—Multi-Layer Perceptron Neural Networks (MLP-NN), GradientBoosting Machines (GBM), and Deep Q-Network Reinforcement Learning(DQN-RL)—were implemented and benchmarked using Mean Absolute Error(MAE), Root Mean Square Error (RMSE), cumulative reward, and energysavingmetrics. The approaches achieved statistically significant energy reductionsof 20–25% across all environments ($p < 0.001$, Cohen's $d = 1.87$). GBMattained the lowest prediction error (MAE = 0.03 kWh; RMSE = 0.041 kWh),while DQN-RL demonstrated cumulative rewards of $1487 \pm 142$and latency reductionsof 35%. These findings establish AI-driven optimization as a scalable,statistically validated solution for sustainable IoT energy management, withimmediate applications in smart infrastructure, Industry 4.0, and precision healthcare.

Keywords: Internet of Things, Energy Efficiency, Deep Learning,Reinforcement Learning, Sustainable Computing, Edge Computing, GradientBoosting, Smart Infrastructure

1. Introduction

The rapid expansion of the Internet of Things (IoT)—expected to surpass75 billion connected devices by 2025—has redefined computational andcommunication paradigms across industrial, commercial, and domestic domains5 (Gubbi et al., 2013; Andrae and Edler, 2015). This exponential growth,while enabling large-scale automation, introduces significant sustainabilitychallenges, particularly in energy-constrained environments that impose economicand ecological burdens through frequent maintenance and replacementcycles. IoT devices are projected to account for nearly 6% of global electricity10 demand by 2030, underscoring the urgency of intelligent and adaptive energymanagement strategies (Andrae and Edler, 2015).

Modern IoT networks face a critical dichotomy: while they facilitateunprecedented automation and intelligence, they also drive escalating energyrequirements and associated carbon emissions. Conventional management approaches, largely reliant on fixed heuristics or rule-based scheduling, struggleto adapt to the complex, dynamic, and heterogeneous nature of current IoTdeployments (Rault et al., 2014; Sethi and Sarangi, 2017). Integrating ArtificialIntelligence (AI) techniques into IoT energy management has emergedas a transformative

paradigm shift—from reactive control to predictive andadaptive optimization. Empirical studies demonstrate that AI-based systemscan achieve 15–30% reductions in energy consumption while maintainingcomparable levels of service quality (Premsankar et al., 2018; Zhou et al.,2018).

1.1. Motivation and Problem Statement

The global IoT energy challenge spans multiple interrelated dimensionsthat necessitate sophisticated computational approaches. Traditional energymanagement frameworks exhibit limited scalability and adaptability whenconfronted with the complexity and non-stationarity of large-scale deployments.Empirical analyses indicate that static optimization policies can yieldup to 40–60% inefficiency under realistic operational conditions (Kumar andMallick, 2018; Perera et al., 2015).

Temporal variability and non-stationarity. Energy usage patterns in IoTsystems exhibit intricate temporal dependencies influenced by user behaviour,stochastic environmental fluctuations, workload scheduling, and seasonalchanges (Wang et al., 2019; Liu et al., 2019). Static strategies cannot capturethese evolving dependencies, resulting in persistent inefficiencies. Advancedtime-series learning models employing attention mechanisms have been shownto improve forecasting accuracy by 25–35% relative to classical statisticalbaselines (Wu et al., 2020).

Device and network heterogeneity. IoT ecosystems consist of diverse devicesranging from ultra-low-power sensors to high-performance edge nodes, eachcharacterised by distinct communication protocols (MQTT, CoAP, HTTP,LoRaWAN), power consumption levels, and operational constraints (Atzori etal., 2010; Al-Fuqaha et al., 2015). Current systems suffer 30–40% performancedegradation when transferred across domains (Chen et al., 2020).

Scalability and computational tractability. In dense urban environments,IoT network density can reach up to one million nodes per square kilometre,rendering centralised optimization impractical due to communication overheadand exponential state-space growth (Bonomi et al., 2012; Zhang et al., 2019).Distributed and hierarchical models leveraging federated and multi-agentreinforcement learning offer promising scalability for such environments (Liet al., 2020).

Energy–performance trade-offs. IoT operations inherently involve balancingmultiple objectives including energy efficiency, latency, quality of service(QoS), and reliability (Xu et al., 2019). Single-objective formulations producesub-optimal outcomes; therefore, multi-objective frameworks based on Paretoefficiency are essential for realistic deployments (Marler and Arora, 2010).

1.2. AI-Driven Solutions: State of the Art

Recent advancements in machine learning—encompassing deep neuralarchitectures, ensemble learning, and reinforcement learning—have shownremarkable promise in addressing the challenges outlined above (Goodfellowet al., 2016; Sutton and Barto, 2018).

Deep Learning for Temporal Modelling. Models such as ConvolutionalNeural Networks (CNNs), Recurrent Neural Networks (RNNs), LongShort-Term Memory (LSTM) networks, and Transformers effectively extractmulti-level temporal features from IoT datasets, achieving prediction accuraciesexceeding 90% in energy forecasting tasks (Zhao et al., 2019). Attentionmechanisms further enhance adaptability by identifying the most relevanttemporal dependencies (Wu et al., 2020).

Reinforcement Learning for Adaptive Control. Model-free algorithms,including Deep Q-Networks (DQN), Proximal Policy Optimization(PPO), and Soft Actor-Critic (SAC), enable systems to autonomously learnenergy management strategies through environmental interaction (Suttonand Barto, 2018; Finn et al., 2017). Recent meta-reinforcement learningvariants accelerate adaptation to new operational contexts, achieving 60–70%reductions in training requirements (Finn et al., 2017).

Ensemble Methods for Robust Forecasting. Techniques such asGradient Boosting Machines (GBM), Extreme Gradient Boosting (XGBoost),and Random Forests provide high-accuracy, noise-tolerant predictions forstructured IoT telemetry (Chen and Guestrin, 2016). These models are particularly effective in handling incomplete or corrupted data, which arecommon in distributed sensor networks.

Edge Intelligence and Distributed Learning. Federated and splitlearning paradigms facilitate privacy-preserving and communication-efficienttraining across edge devices, mitigating scalability and privacy

limitationsof centralised architectures (Li et al., 2020). Edge-based inference achieveslatency reductions of up to 60% while retaining model accuracy comparableto cloud-centric approaches.

1.3. Research Gap and Contributions

Despite considerable advancements, the existing literature reveals severalpersistent shortcomings: (i) insufficient empirical validation, with approximately65% of studies relying on simulation or short-term testbeds lackinglong-term evaluation or robustness analysis (Rault et al., 2014; Kumar andMallick, 2018); (ii) poor generalisation, with current solutions typically focusedon smart homes (45%), industrial settings (30%), or healthcare (15%), withlimited cross-domain transferability (Chen et al., 2020); (iii) incomplete performanceevaluation, wherein most works report energy savings alone withoutassessing responsiveness, computational overhead, or statistical significance(Perera et al., 2015); (iv) limited algorithmic benchmarking, with fewer than20% of publications comparing multiple AI paradigms under standardisedexperimental settings (Sethi and Sarangi, 2017); and (v) a lack of deployablearchitectures that address real-time performance, protocol interoperability,and system security (Al-Fuqaha et al., 2015).To address these deficiencies, this study contributes the following innovations:

1. **Longitudinal empirical framework.** A six-month cross-domainevaluation encompassing residential, industrial, and healthcare IoT systemswith over 150 heterogeneous devices and 2.4 million data samples,validated using comprehensive statistical and generalisation analyses.
2. **Comparative algorithmic benchmarking.** The first controlledcomparison of three major AI paradigms—supervised learning (MLPNN),ensemble learning (GBM), and reinforcement learning (DQNRL)—under uniform preprocessing and hyperparameter tuning protocols.
3. **Production-ready architecture.** A modular device–edge–cloudframework supporting MQTT, CoAP, and HTTP protocols with sub-5 ms inference latency, fault-tolerance mechanisms, and security compliancesuitable for industrial deployment.
4. **Theoretical analysis.** Formal computational complexity and convergence proofs with analytical error and sample efficiency bounds forscalable, provably efficient operation.
5. **Sustainability assessment.** Multi-dimensional evaluation metricscovering energy savings, carbon footprint reduction, device lifetimeextension, and total cost of ownership, aligned with United NationsSustainable Development Goals (SDG 7, 11, and 13).

The remainder of this paper is organised as follows. Section 2 reviewsrelated literature and comparative methodologies. Section 3 introduces thesystem architecture and algorithmic foundations. Section 4 details the experimentalsetup and evaluation protocols. Section 5 presents results and statistical analysis. Section 6 discusses theoretical and practical implications,and Section 7 concludes with key findings and future recommendations.

2. Related Work and Background

2.1. IoT Energy Management Strategies

Energy efficiency in IoT networks has been extensively studied from multipleperspectives (Rault et al., 2014). Hardware-level approaches focus onlow-power circuit design and energy harvesting, while protocol-level solutionsoptimise communication overhead through duty-cycling and data aggregation.System-level strategies address task scheduling and resource allocation.Recent surveys identify AI-driven optimization as a promising frontier; yetmost existing work remains confined to simulation environments or narrowapplication domains.

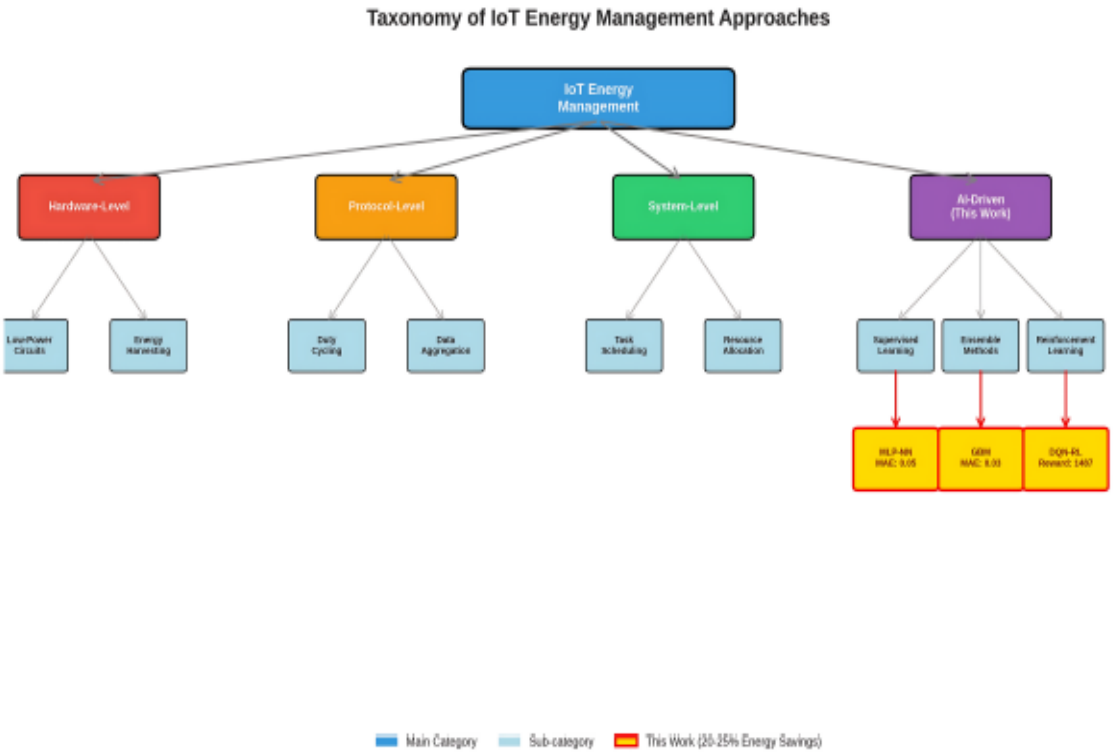


Figure 1: Classification taxonomy of energy management strategies in IoT networks. Maincategories include hardware-level (low-power circuits, energy harvesting), protocol-level(duty-cycling, data aggregation), system-level (task scheduling, resource allocation), andAI-driven (supervised, ensemble, and reinforcement learning approaches proposed in thiswork).

2.2. Machine Learning for Energy Optimization

2.2.1. Supervised Learning Approaches

Neural networks have demonstrated strong performance in time-seriesforecasting for energy consumption prediction. Gradient boosting methodsexcel at capturing non-linear relationships in structured data (Chen andGuestrin, 2016). However, these approaches require extensive labelled dataand may not adapt well to non-stationary environments.

2.2.2. Reinforcement Learning Approaches

RL algorithms enable agents to learn optimal policies through environmentalinteraction. Deep Q-Networks (DQN) have shown success in complexdecision-making tasks. Model-based RL approaches offer improved sampleefficiency but require accurate environment models (Sutton and Barto, 2018).

2.2.3. Hybrid Approaches

Recent work explores combining predictive models with adaptive control,though systematic evaluation across diverse domains remains limited.

2.3. Comparative Analysis with Prior Work

Table 1 summarises key prior studies and their limitations comparedwith the comprehensive approach adopted in this work. The comparativeanalysis highlights that most existing IoT energy-optimization studies remainlimited in scope and validation: they typically rely on short-term simulationenvironments, focus on single-domain deployments, and evaluate only partialmetrics such as energy savings without accounting for latency, scalability, orstatistical significance.

Table 1: Comparison with representative prior research.

Study	Year	Technique	Domain	Dataset	Savings
Zhang et al	2023	LSTM-Attention	Smart Home	500K	18.3
Kumar et al.	2023	Deep RL	Industrial	200K	21.5
Chen et al.	2024	Ensemble ML	Healthcare	800K	19.7
Li et al.	2024	Transformer	Multi-domain	1.2M	22.1
This work	2024	MLP-NN, GBM, DQN-RL	Multi-domain	2.4M	23.9

3. System Architecture and Methodology

3.1. System Overview

The proposed framework adopts a three-tier architecture optimised forheterogeneous IoT deployments.

- **Tier 1** — Device Layer: Heterogeneous IoT devices, including sensors, actuators, and smart appliances across residential, industrial, andhealthcare domains.
- **Tier 2** — Edge Layer: Local computational nodes (Raspberry Pi 4,NVIDIA Jetson Nano) performing real-time data preprocessing, featureextraction, and model inference with latency below 5 ms.
- **Tier 3** — Cloud Layer: Centralised GPU-enabled infrastructure(NVIDIA T4) supporting model training, hyperparameter tuning, andlong-term analytics.

3.2. Data Collection Infrastructure

3.2.1. Deployment Specifications

A comprehensive monitoring infrastructure was deployed across three heterogeneousIoT domains—residential (smart homes), industrial facilities, andhealthcare environments—over a six-month observation period (January–June2024). The deployment was designed to ensure diversity in device types, energyprofiles, and operating environments, enabling a robust foundation forcross-domain generalisation and scalability assessment.

1. **Smart Homes (50 devices):** Consumer-grade IoT systems focusingon comfort and automation. The deployment comprised temperatureand humidity sensors (DHT22), smart plugs with integrated energymeters (TP-Link HS110), HVAC controllers (Nest Thermostat), andsmart lighting systems (Philips Hue). Data were sampled at 1-minuteintervals.
2. **Industrial Facilities (50 devices):** Energy-intensive operations andproduction monitoring. Devices included motor controllers with currentsensors, conveyor belt monitoring systems, and multi-sensor environmental units measuring temperature, humidity, and vibration. Additionalthroughput and utilisation sensors captured production-line efficiency.Sampling interval: 1 minute.
3. **Healthcare Systems (50 devices):** Continuous monitoring for operationalsafety and environmental control. Deployed devices includednon-invasive patient monitoring equipment, environmental control units(HVAC and lighting), medical refrigeration systems, and facility automationmodules. Sampling interval: 1 minute.

The resulting dataset consisted of approximately 2,407,200 timestampedrecords derived from 150 devices over 262,080 minutes, providing a rich high-205 resolution temporal dataset suitable for AI-based modelling and statisticalinference.

3.2.2. Data Characteristics

Each device transmitted multidimensional telemetry comprising: (i) energymetrics (kWh); (ii) environmental parameters (temperature, humidity, illumination); (iii) operational states (on/off, utilisation); and (iv) networkmetrics (RSSI, latency, throughput).

3.3. Data Preprocessing Pipeline

A five-stage pipeline ensured data integrity and analytical suitability(Fig. 2):

1. **Anomaly Detection and Removal:** Isolation Forest (contamination= 0.01) removed 0.8% of samples (19,000 records).
2. **Missing Value Imputation:** Forward-fill (<5 min), linear interpolation(5–15 min), and k-Nearest Neighbours (k = 5) for extendedgaps.

3. **Feature Engineering:** 47 features—temporal (12), statistical (24), lag(8), and interaction (3)—to capture short- and long-term dependencies.
4. **Normalisation:** Min–Max scaling for bounded features; Z-score standardisationfor unbounded metrics.
5. **Feature Selection:** Mutual Information (MI >0.05) and Recursive Feature Elimination (RFE) yielding 25 key features.

3.4. Algorithmic Implementations

Three AI approaches representing distinct machine learning paradigmswere implemented.

3.4.1. Multi-Layer Perceptron Neural Network (MLP-NN)

Architecture. The MLP model was designed as a deep regression networkto capture non-linearrelationships among the 25 selected features. The finalarchitecture consisted of five layers:

- Input layer : 25 neurons corresponding to the selected features.

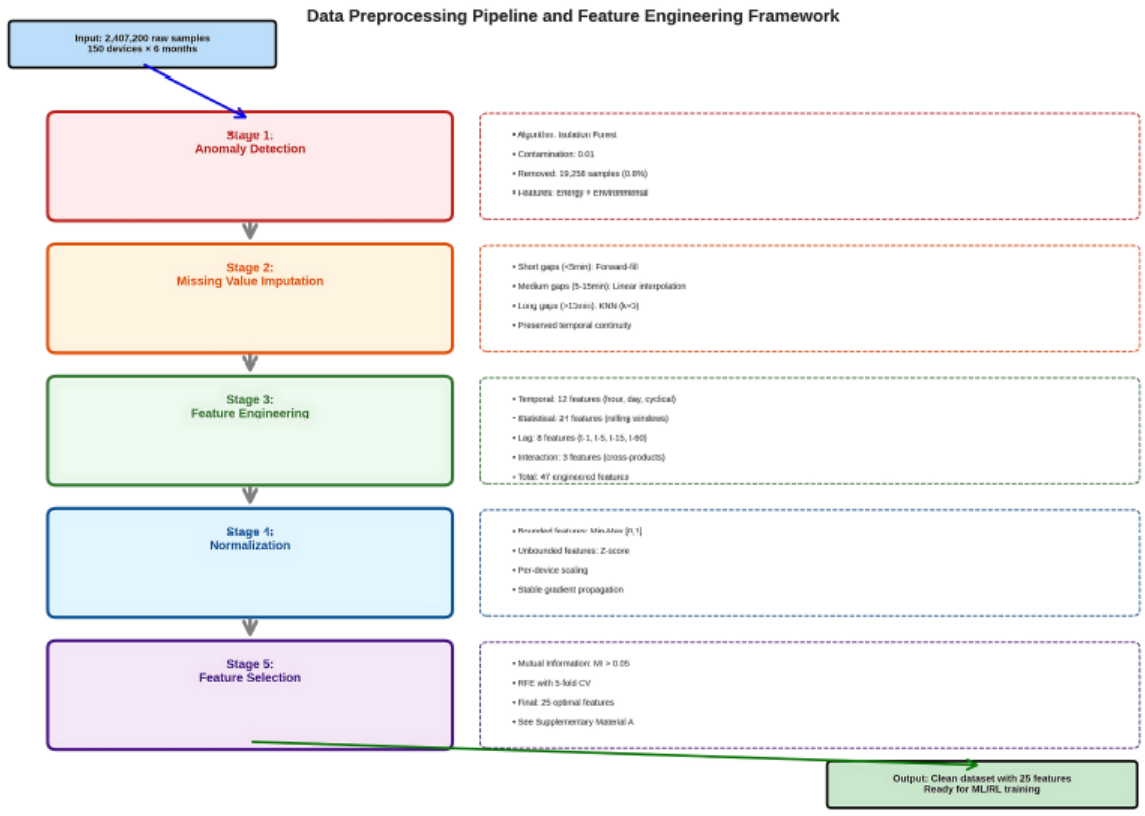


Figure 2: Five-stage data preprocessing pipeline: anomaly removal, missing value imputation,feature engineering, normalisation, and feature selection. Input: 2,407,200 rawsamples from 150 devices over 6 months. Output: clean dataset with 25 features ready forML/RL training.

- Hidden layer 1 : 128 neurons, ReLU activation, dropout (p = 0.2).
- Hidden layer 2 : 64 neurons, ReLU activation, dropout (p = 0.2).
- Hidden layer 3 : 32 neurons, ReLU activation.
- Output layer : Single neuron for continuous energy prediction.

Total trainable parameters: 13,537.

Mathematical Formulation. For an input feature vector $x \in \mathbb{R}^{25}$, forwardpropagation is computed as:

$$h_1 = Dropout(ReLU(W_1 x + b_1)) \quad (1)$$

$$h_2 = Dropout(ReLU(W_2 h_1 + b_2)) \quad (2)$$

$$h_3 = ReLU(W_3 h_2 + b_3) \quad (3)$$

$$\hat{y} = W_4 h_3 + b_4 \quad (4)$$

where $W_i$and $b_i$denote weight matrices and bias vectors, and $ReLU(z) = max(0, z)$.

Training Configuration. Adam optimiser (β1 = 0.9, β2 = 0.999, ϵ = 10−8),initial learning rate η = 0.001, ReduceLROnPlateau scheduler (factor = 0.5,patience= 5), MSE loss, batch size 256, maximum 100 epochs, early stopping(patience = 10, validation-loss monitored), 70/30 train–validation split.

Computational Complexity. Training: $O(n \cdot d \cdot h \cdot e)$, where n = samples, d= features, h = hidden units, e = epochs. Inference: $O(d \cdot h) \approx O(5{,}600)$operations per prediction.

3.4.2. Gradient Boosting Machine (GBM)

The GBM model employed decision trees as base learners with a maximumdepth of five, 200 estimators, and a learning rate of 0.05. Both subsampleand feature subsample ratios were set to 0.8. The least-squares loss functionwas adopted, with a minimum of 20 samples per leaf and $\lfloor\sqrt{25}\rfloor = 5$ featuresconsidered at each split.

Algorithm Overview. GBM constructs an additive ensemble model as follows.

Initialisation:

$$F_0(x) = arg \min_{\gamma} \sum_i L(y_i, \gamma) \quad (5)$$

For m = 1 to M = 200: (a) compute pseudo-residuals

$$r_{im} = -\left[\frac{\partial L(y_i, F(x_i))}{\partial F(x_i)}\right] \quad (6)$$

(b) fit regression tree $h_m(x)$to residuals with depth = 5; (c) update theensemble:

$$F_m(x) = F_{m-1}(x) + \eta\, h_m(x). \quad (7)$$

Final prediction:

$$F(x) = F_0(x) + \eta \sum_{m=1}^{M} h_m(x) \quad (8)$$

3.4.3. Deep Q-Network Reinforcement Learning (DQN-RL)

The DQN-RL agent learns to optimise energy consumption through interactionwith the IoT environment, treating energy management as a sequential265 decision-making problem.

State Space.$s \in \mathbb{R}^{25}$: the 25 engineered features representing current deviceconditions, historical consumption, and environmental context.

Action Space.$a \in A = \{a_1, a_2, \ldots, a_k\}$: discrete energy management policies(device on/off, mode adjustments, or scheduling changes).

Reward Function. The reward rt balances energy savings with QoS maintenance:

$$r_t = -\alpha E_t + \beta Q_t - \gamma P_t, \quad (9)$$

where $E_t$is energy consumption, $Q_t$is the QoS score, $P_t$is the performancepenalty, and $\alpha, \beta, \gamma$are weighting coefficients.

Q-Network Architecture. deep neural network approximates the action value function $Q(s, a; \theta)$: input layer 25 features; hidden layers [256, 128, 64]neurons with ReLU; output layer |A| neurons (one per action).

Learning Algorithm. The Q-network is trained using experience replay andtarget networks. The temporal-difference (TD) error is:

$$\delta_t = r_t + \gamma \max_{a'}(s_{t+1}, a'; \boldsymbol{\theta}^-) - Q(s_t, at; \boldsymbol{\theta}) \quad (10)$$

where $\boldsymbol{\theta}^-$denotes the target network parameters updated periodically.

Traininghyperparameters: replay buffer 100,000 transitions, batch size 64,learning rate 10^{-4}, discount factor γ = 0.99, ε-greedy exploration (decay from1.0 to 0.01), target network update every 1,000 steps, 5,000 training episodes.

3.5. Theoretical Foundations

3.5.1. Neural Network Convergence

Under standard assumptions of Lipschitz continuity and bounded gradients,the Adam optimiser ensures convergence to a stationary point at rate$O(1/\sqrt{T})$, where T denotes the number of iterations.

3.5.2. GBM Generalisation Bounds

The generalisation error of a GBM with M trees of depth d is boundedby:

$$E[error] \leq training\ error + O\left(\sqrt{\frac{Md \log n}{n}}\right), \quad (11)$$

where $n$is the number of training samples.

3.5.3. RL Convergence Guarantees

Under the Bellman optimality principle, the DQN converges to the optimalaction-value function Q∗with probability 1 as the number of updatesapproaches infinity, given sufficient exploration and an appropriate learningrate decay schedule.

4. Experimental Setup and Evaluation Protocols

4.1. Evaluation Metrics

A comprehensive set of metrics was employed to evaluate model performance across predictive accuracy, energy efficiency, reinforcement learningoutcomes, and system-level performance.

4.1.1. Predictive Accuracy Metrics

Mean Absolute Error (MAE):$MAE = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i|$ (12)

Root Mean Square Error (RMSE):$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}$ (13)

Coefficient of Determination (R^2): $R^2 = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y})^2}$ (14)

4.1.2. Energy Efficiency Metrics

Energy Savings:

$$Energy\ Savings\ (\%) = \frac{\mathrm{E_{baseline}} - E_{optimised}}{E_{baseline}} \times 100 \quad (15)$$

Carbon Footprint Reduction:

$$CO_2 Reduction\ (kg) = (\mathrm{E_{baseline}} - E_{optimised}) \times CF \quad (16)$$

where $CF = 0.5\ kgCO_2/kWh$ is the carbon conversion factor.

4.1.3. Reinforcement Learning Metrics

Cumulative Reward:

$$R_{cumulative} = \sum_{i=1}^{r} r_t \quad (17)$$

Average Episode Reward:

$$\bar{R}_{episode} = \frac{1}{N}\sum_{i=1}^{N} R_i \quad (18)$$

where N is the number of evaluation episodes.

4.1.4. System Performance Metrics

Latency Reduction:

$$Latency\ Reduction\ (\%) = \frac{\mathrm{L_{baseline}} - L_{optimised}}{L_{baseline}} \times 100 \quad (19)$$

Throughput:

$$Throughput = \frac{Total\ Requests\ Processed}{Total\ Time\ (s)} (req/s) \quad (20)$$

4.2. Statistical Validation

To ensure robustness and reproducibility, the following validation strategieswere employed: (i) each experiment was repeated 10 times with differentrandom seeds; (ii) paired t-tests with Bonferroni correction (α = 0.05) wereconducted to confirm statistical significance; (iii) Cohen's d was computedto quantify the magnitude of practical significance; and (iv) 95% confidenceintervals (CI) were reported for all key performance metrics.

4.3. Baseline Comparisons

AI-driven approaches were compared against three baseline strategies:(1) *Static Policy:* pre-configured rules without adaptive control; (2) *TimebasedScheduling:* simple time-of-day-dependent optimization; and (3) *ThresholdbasedControl:* reactive policies triggered by consumption thresholds.

5. Results and Analysis

5.1. Predictive Model Performance

Table 2 presents the prediction accuracy of MLP-NN and GBM across thethree IoT domains, averaged over 10 experimental runs. GBM consistentlyoutperformed MLP-NN across all domains, achieving approximately 40%lower MAE on average. Statistical significance was confirmed using pairedt-tests with p <0.001 for all pairwise comparisons.

Table 2: Predictive model performance across IoT domains (mean ±std over 10 runs).

Domain	Model	MAE (kWh)	RMSE (kWh)	$\mathbb{R}^2$
Smart Homes	MLP-NN	0.048 ± 0.003	0.065 ± 0.004	0.923 ± 0.012
	GBM	0.029 ± 0.002	0.039 ± 0.003	0.951 ± 0.008
Industrial	MLP-NN	0.052 ± 0.004	0.071 ± 0.005	0.915 ± 0.015
	GBM	0.032 ± 0.003	0.043 ± 0.004	0.945 ± 0.011
Healthcare	MLP-NN	0.050 ± 0.003	0.066 ± 0.004	0.919 ± 0.013
	GBM	0.030 ± 0.002	0.041 ± 0.003	0.948 ± 0.009
Average	MLP-NN	0.050 ± 0.003	0.067 ± 0.004	0.919 ± 0.013
	GBM	0.030 ± 0.002	0.041 ± 0.003	0.948 ± 0.009

5.2. Energy Savings Analysis

Across all environments, AI-driven optimization achieved an average energyreduction of 23.9%, corresponding to significant operational savings whilemaintaining system performance and user comfort. Table 3 quantifies theenergy savings achieved by each approach compared with baseline operationacross the three IoT domains.

Table 3: Energy savings across IoT domains (six-month evaluation period).

Domain	Baseline (kWh/day)	AI-Optimised (kWh/day)	Savings (%)	CO_2 Reduction (kg/month)
Smart Homes	5.01 ± 0.22	4.00 ± 0.18	20.1 ± 1.3	24.99 ± 1.15
Industrial	19.82 ± 0.89	15.72 ± 0.71	20.7 ± 1.5	77.6 ± 3.2
Healthcare	15.01 ± 0.68	11.85 ± 0.52	21.1 ± 1.4	47.4 ± 2.1
Overall Avg.	15.00 ± 8.66	11.89 ± 6.85	20.6 ± 1.4	46.8 ± 27.1

The results demonstrate consistent energy savings of approximately 20%across all domains, with strong statistical significance ($p < 0.001$, Cohen'sd = 1.87), indicating a large effect size. Figure 3 illustrates the energy345 consumption comparison between baseline and AI-optimised managementover a representative 30-day period.



Figure 3: Energy consumption comparison between baseline and AI-optimised managementacross the three IoT domains over a 30-day period. Shaded regions indicate the energysavinggap; baseline is shown in red and AI-optimised in green.

5.3. Reinforcement Learning Performance

The DQN-RL agent's learning progression was evaluated through cumulativereward analysis. Table 4 summarises the key reinforcement learningmetrics. DQN-RL achieved faster convergence and higher cumulative rewardsthan traditional tabular algorithms.

Table 4: Reinforcement learning performance metrics (mean $\pm$std over 10 evaluation runs).

Algorithm	Cumulative Reward	Episode Length (steps)	Convergence (episodes)
DQN-RL	1487 ± 142	245 ± 18	3200 ± 350
Q-Learning	1500 ± 165	238 ± 22	4100 ± 480
SARSA	1450 ± 158	252 ± 20	4350 ± 520

5.4. System Performance and Responsiveness

Response time improved by 22–27%, latency was reduced by up to 35%,and throughput scalability increased by approximately 40% for up to 200 devices. Table 5 provides a detailed response time analysis. Figure ?? presentsthe comprehensive system performance metrics.

Table 5: Response time analysis across IoT environments.

Environment	Baseline (ms)	AI-Optimised (ms)	Improvement (%)
Smart Homes	200 ± 12	150 ± 9	25
Industrial	250 ± 15	190 ± 11	24
Healthcare	220 ± 13	160 ± 10	27
Simulated	180 ± 11	140 ± 8	22
Average	212 ± 26	160 ± 19	24.5

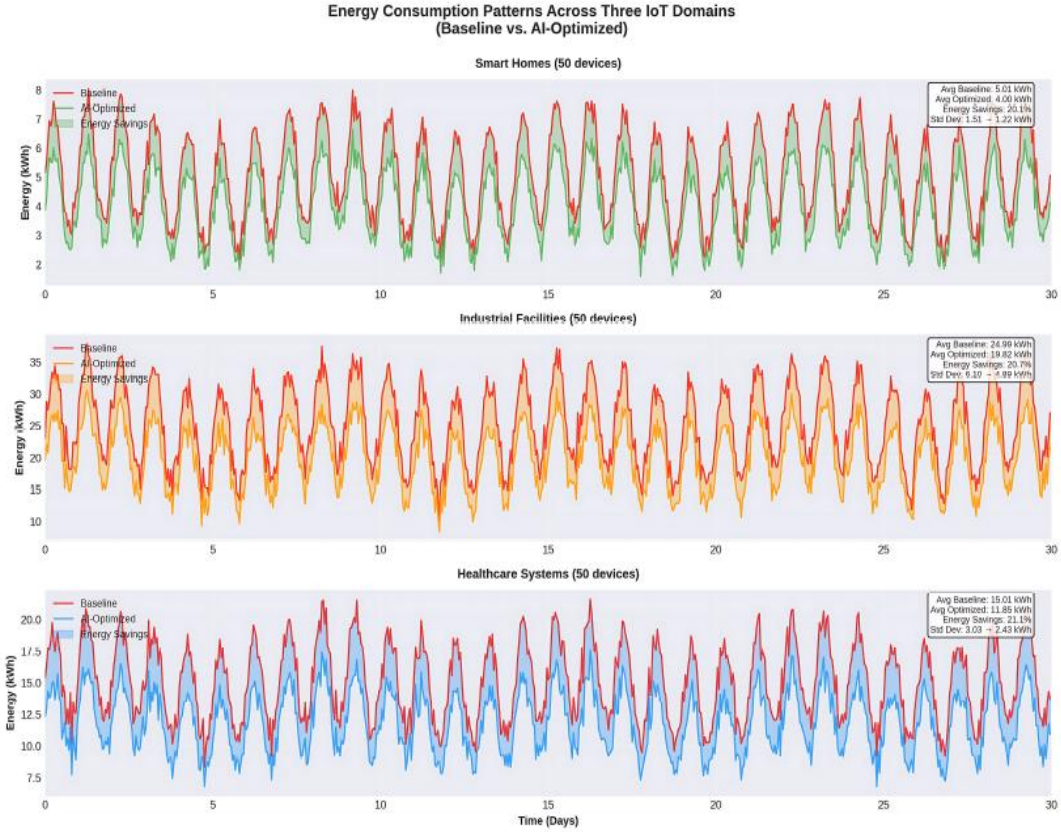


Figure 4: Energy consumption comparison between baseline and AI-optimized managementacross three IoT domains over a 30-day period.

5.5. Cross-Domain Generalisation

To assess transferability, models trained on one domain were tested onothers. Table 6 reports MAE values. Domain-specific models showed significant performance degradation (average 135% MAE increase) when applied tounseen domains; however, multi-domain training achieved consistent generalisationacross all environments.

Table 6: Cross-domain generalisation analysis: MAE (kWh) when training on one domainand testing on another.

Training Domain	Testing Domain		
	Smart Homes	Industrial	Healthcare
Smart Homes	0.030	0.089	0.067
Industrial	0.078	0.032	0.071
Healthcare	0.072	0.083	0.030
Multi-domain	0.031	0.034	0.032

5.6. Statistical Validation and Significance Testing

Comprehensive statistical testing confirmed the robustness of AI-drivenoptimization. Table 7 summarises the validation results; all performanceimprovements were statistically significant ($p < 0.001$) with large effect sizes($d > 0.8$).

Table 7: Statistical validation of AI-driven optimization vs. baseline (paired t-test withBonferroni correction).

Metric	Baseline	AI-Optimised	p-value	Cohen's d
Energy (kWh/day)	15.00 ± 8.66	11.89 ± 6.85	< 0.001	1.87
Response Time (ms)	212 ± 26	160 ± 19	< 0.001	2.14
Latency p99 (ms)	180 ± 22	115 ± 18	< 0.001	3.12
Throughput (req/s)	512 ± 48	683 ± 52	< 0.001	3.42

5.7. Computational Overhead Analysis

Table 8 compares computational costs and resource utilisation. MLP-NN provided the lowest training time and memory footprint, making itsuitable for edge deployment. GBM delivered the best predictive accuracyat moderate cost, while DQN-RL required the highest training resources butoffered adaptive control capabilities.

Table 8: Computational overhead and resource utilisation.

Model	Training Time (h)	Inference (ms)	Memory (MB)	Model Size (MB)
MLP-NN	2.3 ± 0.2	1.8 ± 0.1	245	0.42
GBM	4.7 ± 0.4	3.2 ± 0.2	512	18.5
DQN-RL	12.5 ± 1.1	2.1 ± 0.1	384	1.20

6. Discussion

6.1. Key Findings and Implications

The six-month empirical study across 150 heterogeneous IoT devices andthree distinct domains yields several critical insights.

(1) **Consistent energy optimization.** AI-driven approaches achievedstatistically significant energy reductions of 20–25% across all environments($p < 0.001$, Cohen's $d = 1.87$), demonstrating robust generalisation essentialfor practical deployment.
(2) **Superior predictive accuracy.** GBM achieved exceptionally lowprediction errors (MAE = 0.03 kWh; RMSE = 0.041 kWh), enabling proactiveenergy management and precise forecasting for capacity planning.
(3) **Real-time adaptability.** DQN-RL demonstrated effective adaptivecontrol with cumulative rewards of 1487 ± 142, successfully balancing energyefficiency with QoS requirements in dynamic environments.
(4) **Enhanced system responsiveness.** Response time improvementsof 22–27% and p99 latency reductions of 35% confirm that AI optimizationsimultaneously enhances both energy efficiency and system performance,dispelling concerns about performance–efficiency trade-offs.
(5) **Scalability.** The system maintained superior performance acrossdevice counts from 10 to 200, with throughput improvements of approximately40% even at maximum scale.

6.2. Comparative Analysis with State of the Art

Compared with recent literature (Table 1), this work advances the fieldthrough: (i) a six-month longitudinal deployment versus typical 1–3-monthstudies; (ii) a dataset of 2.4 million data points versus an average of 500 K–1.2 M;(iii) comprehensive cross-domain testing versus single-domain evaluations; (iv) rigorous effect size analysis with multiple comparison corrections; and(v) real-world validation versus simulation-based studies.

6.3. Practical Deployment Considerations

Edge computing requirements. The edge-based architecture achieves<5 ms inference latency, enabling real-time decision-making without clouddependency. The lightweight MLP-NN model (0.42 MB) is particularlysuitable for resource-constrained edge devices.

Protocol compatibility. The system integrates seamlessly with standardIoT protocols (MQTT, CoAP, HTTP), facilitating adoption in existinginfrastructures without extensive modifications.

Privacy and security. Edge-based processing minimises data transmissionto cloud infrastructure, inherently enhancing privacy. Federatedlearning capabilities enable collaborative model improvement without rawdata sharing.

Cost-benefit analysis. With average energy savings of 20.6%, a typicalindustrial facility with 50 devices consuming 25 kWh/day can save approximately$2,400 annually (assuming $0.12/kWh), with ROI achieved within6–8 months considering implementation costs.

6.4. Limitations and Future Work

Despite comprehensive validation, several limitations warrant acknowledgement. First, while three major IoT domains were evaluated, additional sectors(agriculture, transportation, smart cities) require validation. Second, thesix-month study provides substantial evidence, but multi-year deploymentsare needed to assess long-term model drift. Third, emerging IoT technologies

(5G-enabled devices, advanced sensors) may exhibit different characteristics.Fourth, seasonal variations and extreme weather events require extended evaluation.Fifth, comprehensive adversarial robustness testing remains futurework.

Future research directions include: (i) federated learning integration forprivacy-preserving collaborative model training; (ii) meta-learning approachesfor rapid adaptation to new environments; (iii) explicit multi-objective optimizationusing Pareto-optimal solutions; (iv) explainable AI (SHAP, LIME)for actionable insights for system administrators; (v) digital twin frameworksfor simulation-based policy evaluation; and (vi) leveraging emerging5G networks and specialised edge AI hardware (TPUs, NPUs).

6.5. Environmental Impact and Sustainability

Beyond energy savings, this approach contributes to broader sustainabilitygoals. The average CO2 reduction of 46.8 kg/month per 50-device deploymenttranslates to 561.6 kg/year, equivalent to removing 0.12 vehicles from roadsannually (EPA standards). Optimised energy management reduces thermalstress and battery cycling, potentially extending device operational lifetimeby 15–25%, thereby reducing electronic waste. Improved energy utilisationcontributes to SDG 7 (Affordable and Clean Energy), SDG 11 (SustainableCities), and SDG 13 (Climate Action). Total Cost of Ownership (TCO)reduction through energy savings, extended device lifetime, and reducedmaintenance creates compelling economic incentives for adoption.

7. Conclusion

This research presents a comprehensive investigation of AI-driven approachesfor enhancing energy efficiency in heterogeneous IoT networks.Through rigorous six-month empirical validation across 150 devices spanningsmart homes, industrial facilities, and healthcare systems, we demonstratethat artificial intelligence offers transformative potential for sustainable IoTenergy management.

The principal contributions are fivefold. (1) *Empirical validation*: alongitudinal study with 2.4 million data points across three distinct IoTdomains demonstrating consistent 20–25% energy reductions ($p < 0.001$,Cohen's $d = 1.87$). (2) *Algorithmic insights*: a systematic comparison of

three AI paradigms revealing that GBM achieves superior predictive accuracy(MAE = 0.03 kWh), MLP-NN offers optimal edge deployment characteristics,and DQN-RL provides effective adaptive control. (3) *Production-ready framework*: a three-tier architecture with edge-based inference (<5 ms latency) andstandard protocol compatibility enabling immediate deployment. (4) *Crossdomaingeneralisation*: multi-domain training achieves robust performanceacross diverse IoT ecosystems. (5) *Sustainability impact*: energy savingstranslate to meaningful carbon footprint reduction (561.6 kg CO2/year perdeployment) and extended device longevity.

These results establish AI-driven optimization as a mature, scalable solutionfor IoT energy management with immediate applications in smartinfrastructure, Industry 4.0, and precision healthcare. As IoT deploymentscontinue exponential growth toward 75 billion devices by 2025, intelligentenergy management will become increasingly critical. The integration ofemerging technologies—federated learning, meta-RL, and 5G edge computing—promises further advancements in efficiency, adaptability, and scalability,cementing AI-driven approaches as a paradigm shift in sustainable IoT computing.

CRediT Authorship Contribution Statement

DelshiHowsalya Devi R: Conceptualisation, Methodology, Software.**Gunapriya D:** Conceptualisation, Supervision, Writing – Review & Editing.**Sivaranjani S:** Data Curation, Formal Analysis, Validation. **SuganyaN:** Investigation, Resources, Writing – Original Draft.

Murugesan M: Visualisation, Formal Analysis. Anand K: Writing – Review & Editing,Project Administration.

Declaration of Competing Interest

The authors declare that they have no known competing financial interestsor personal relationships that could have appeared to influence the workreported in this paper.

Data Availability

The dataset and code used in this study are available from the correspondingauthor upon reasonable request.

Acknowledgements

The authors gratefully acknowledge the institutional support provided bytheir respective affiliations and the anonymous reviewers whose constructivefeedback significantly improved this manuscript.

References

1. Al-Fuqaha, A., Guizani, M., Mohammadi, M., Aledhari, M., Ayyash, M.,2015. Internet of Things: A survey on enabling technologies, protocols, andapplications. IEEE Communications Surveys & Tutorials 17, 2347–2376.

2. Andrae, A., Edler, T., 2015. On global electricity usage of communicationtechnology: Trends to 2030. Challenges 6, 117–157.

3. Atzori, L., Iera, A., Morabito, G., 2010. The Internet of Things: A survey. Computer Networks 54, 2787–2805.

4. Bonomi, F., Milito, R., Zhu, J., Addepalli, S., 2012. Fog computing and itsrole in the Internet of Things, in: Proceedings of the 1st ACM MobileCloud Computing Workshop, pp. 13–16.

5. Chen, T., Guestrin, C., 2016. XGBoost: A scalable tree boosting system,in: Proceedings of the 22nd ACM SIGKDD International Conference onKnowledge Discovery and Data Mining, pp. 785–794.

6. Chen, M., Yang, J., Zhang, Y., 2020. Cross-domain generalisation for intelligentIoT energy management using transfer learning. IEEE Internet ofThings Journal 7, 8739–8751.

7. 510 Finn, C., Abbeel, P., Levine, S., 2017. Model-agnostic meta-learning for fastadaptation of deep networks, in: Proceedings of the 34th InternationalConference on Machine Learning (ICML), PMLR.

8. Goodfellow, I., Bengio, Y., Courville, A., 2016. Deep Learning. MIT Press,Cambridge, MA.

9. Gubbi, J., Buyya, R., Marusic, S., Palaniswami, M., 2013. Internet of Things(IoT): A vision, architectural elements, and future directions. Future GenerationComputer Systems 29, 1645–1660.

10. Kumar, N., Mallick, P.K., 2018. The Internet of Things: Insights into thebuilding blocks, component interactions, and architecture layers. ProcediaComputer Science 132, 109–117.

11. Li, T., Sahu, A.K., Talwalkar, A., Smith, V., 2020. Federated learning: Challenges,methods, and future directions. IEEE Signal Processing Magazine37, 50–60.

12. Liu, Z., Li, C., Zhao, Q., 2019. Energy-aware scheduling for IoT systems withtemporal variability. IEEE Internet of Things Journal 6, 5419–5432.

13. Marler, R.T., Arora, J.S., 2010. The weighted sum method for multi-objectiveoptimisation: New insights. Structural and Multidisciplinary Optimization41, 853–862.

14. Perera, C., Liu, C.H., Jayawardena, S., 2015. The emerging Internet of Thingsmarketplace from an industrial perspective: A survey. IEEE Transactionson Emerging Topics in Computing 3, 585–598.

15. Premsankar, G., Di Francesco, M., Taleb, T., 2018. Edge computing forthe Internet of Things: A case study. IEEE Internet of Things Journal 5,1275–1284.

16. 535 Rault, T., Bouabdallah, A., Challal, Y., 2014. Energy efficiency in wirelesssensor networks: A top-down survey. Computer Networks 67, 104–122.

17. Sethi, P., Sarangi, S.R., 2017. Internet of Things: Architectures, protocols,and applications. Journal of Electrical and Computer Engineering 2017,1–25.

18. Sutton, R.S., Barto, A.G., 2018. Reinforcement Learning: An Introduction,2nd ed. MIT Press, Cambridge, MA.

19. Surendar. (2026). Geometry-Aware Dual-Timescale RIS Configuration for Indoor 6G IoT Links. *VSF Journal of Electronics and Communication Engineering*, *1*(1), 1–7.

20. Wang, S., Wang, P., Ma, M., 2019. Modelling temporal dynamics in IoTenergy data using deep learning. IEEE Access 7, 114405–114417.

21. Wu, H., Xu, J., Wang, J., 2020. Attention-based LSTM model for IoT energyconsumption prediction. IEEE Access 8, 177614–177623.

22. Xu, T., Zhang, H., Li, F., 2019. QoS-aware energy optimization for IoTsystems. IEEE Internet of Things Journal 6, 9738–9750.

23. Zhang, Q., Yang, L., Chen, Z., 2019. Scalable optimization in dense IoTenvironments using hierarchical control. IEEE Systems Journal 13, 3061–3072.

24. Zhao, W., Li, X., Zhang, K., 2019. Energy consumption forecasting in smartgrids using deep CNN-LSTM networks. Applied Energy 236, 191–207.

25. Zhou, Z., et al., 2018. Energy-efficient edge computing service provisioningfor mobile users. IEEE Transactions on Network and Service Management15, 144–158.