A Lightweight Edge-Deployable ANN Model for Real-Time Energy Anomaly Detection in IoT-Driven Smart Grids

Sofiane Benabbes*, Wael Aissaoui, Rabah Boucetti
LAMIS Laboratory, Echahid Cheikh Larbi Tebessi University, Tebessa, Algeria
E-mail: sofiane.benabbes@univ-tebessa.dz, wael.aissaoui@gmail.com, boucetti.rabah30@gmail.com
*Corresponding author

Keywords: Artificial Neural Network (ANN), energy anomaly detection, Internet of Things (IoT), smart cities, edge computing, machine learning, energy management systems

Received: October 25, 2025

The rapid expansion of the Internet of Things (IoT) in smart cities has necessitated efficient, real-time energy anomaly detection. However, complex hybrid deep learning models often exceed the computational capacity of Edge devices. This paper proposes a lightweight, 3-layer Artificial Neural Network (ANN) framework designed for Edge deployment. Using the LEAD (Large-scale Energy Anomaly Detection) dataset, we address class imbalance via the Synthetic Minority Over-sampling Technique (SMOTE). Our model achieves 98.4% accuracy, a macro F1-score of 0.93, and an AUC of 0.91. While these metrics are competitive with state-of-the-art hybrid models, our framework provides a significantly lower memory footprint and sub-millisecond inference latency, making it ideal for resource-constrained Edge environments.

Povzetek: Raziskava predlaga lahko umetno nevronsko mrežo za zaznavanje energetskih anomalij v IoT pametnih mestih, ki omogoča visoko natančnost ter hitro in učinkovito delovanje na omejenih robnih napravah.

1 Introduction

The rapid growth of the Internet of Things (IoT) and smart city infrastructures has led to an unprecedented increase in the amount of energy-related data generated by sensors, meters, and connected devices. These data streams are essential for optimizing energy distribution, improving sustainability, and enhancing operational efficiency in urban environments. However, as energy systems become more interconnected and data-intensive, they are increasingly vulnerable to anomalies, which may arise from equipment malfunction, sensor errors, cyberattacks, or abnormal consumption behaviors [1]. Early and accurate detection of such anomalies is therefore critical to ensuring energy efficiency, reliability, and security in smart city ecosystems.

Traditional rule-based or statistical anomaly detection techniques often fail to address the complex, nonlinear, and high-dimensional nature of IoT energy data. In contrast, machine learning and deep learning approaches have shown remarkable potential in modeling dynamic consumption patterns and identifying deviations that are not easily captured by conventional methods. For instance, [2] combined LSTM (Long Short-Term Memory) with Isolation Forest to detect anomalies in IoT energy usage data. [3] proposed a hybrid CNN–Bidirectional LSTM model which achieved precision of 98.7% and recall of 97.9% in detecting abnormal users in a smart grid subsystem. Nonetheless, these models are typically computationally expensive, require large labeled datasets, and are difficult to deploy in real-time Edge Computing environments, limiting their applicability in large-scale urban energy networks.

To address these challenges, this paper proposes a lightweight Artificial Neural Network (ANN) model designed for efficient and accurate detection of energy anomalies in IoT-based smart grids. The model is trained and evaluated on the LEAD (Large-scale Energy Anomaly Detection) dataset, which simulates real-world energy consumption scenarios in diverse environmental and structural contexts. Through rigorous preprocessing—including data cleaning, normalization, and feature engineering—the proposed framework ensures robust input quality and enhances model generalization. The ANN architecture captures the nonlinear dependencies between energy and environmental variables while maintaining low computational complexity, enabling real-time anomaly detection on Edge devices.

The main contributions of this paper are as follows:

1. Development of a scalable ANN-based framework for detecting energy anomalies in IoT-driven smart city environments.
2. Integration of a comprehensive preprocessing pipeline, including advanced feature engineering and normalization, to improve data quality and learning stability.
3. Comparative evaluation with recent state-of-the-art approaches such as [2] and [3], demonstrating competitive or superior performance with 98.4%

accuracy, 0.93 F1-score, and 0.91 AUC, while preserving computational efficiency.

4. Deployment feasibility on Edge Computing architectures, supporting real-time inference and energy-aware processing for large-scale smart city applications.

This research is guided by two primary questions:

- Can a simplified and lightweight ANN architecture achieve anomaly detection performance comparable to complex deep hybrid models?
- What is the quantitative trade-off between model simplicity and inference efficiency when deployed in Edge-computing environments with constrained resources?

The remainder of this paper is organized as follows: Section 2 reviews related works on energy anomaly detection in IoT systems. Section 3 describes the proposed approach and system architecture. Section 4 details the experimental setup, dataset, and preprocessing steps. Section 5 presents and discusses the results, while Section 6 concludes the paper and outlines future research directions.

2 Related works

Research on anomaly detection in Internet of Things (IoT) environments has grown significantly in recent years, driven by the increasing complexity of sensor networks and the critical need for secure and energy-efficient operations in smart cities. Traditional anomaly detection methods have often proven inadequate in handling the massive data streams, resource constraints, and latency requirements inherent to IoT systems. To address these challenges, scholars have explored a wide range of approaches, including lightweight machine learning techniques, deep learning models, edge and fog computing solutions, as well as federated and split learning frameworks. These studies collectively aim to enhance detection accuracy, reduce computational overhead, and improve adaptability across heterogeneous IoT ecosystems. The following review synthesizes these contributions, highlighting key methods, results, and insights that inform the design of reliable and scalable anomaly detection mechanisms for IoT networks.

[2] propose a machine learning framework that combines energy forecasting and unsupervised anomaly detection in IoT networks. Using Long Short-Term Memory (LSTM) networks to predict energy consumption and Isolation Forest to detect anomalies from prediction residuals, their model achieves 98% accuracy. This approach enhances energy efficiency while enabling early threat detection, proving particularly effective for sensitive applications like healthcare.

[4] investigates anomaly detection for IoT cyberattacks in s18mart cities using federated and split learning. The framework balances data privacy with detection performance, enabling collaborative yet privacy-preserving anomaly detection. While the study does not focus specifically on energy anomalies, it offers valuable insights into securing IoT networks against distributed threats in urban environments.

[5] present a real-time anomaly detection framework for smart city IoT sensor data. Their approach integrates unsupervised machine learning with statistical analysis and expert feature engineering to manage large-scale, diverse, and high-velocity data streams. Empirical validation on smart city datasets demonstrates that their model outperforms established anomaly detection techniques, strengthening operational efficiency and urban security.

[6] investigate a deep learning–based anomaly detection system for IoT security in smart cities. Using the IoT-23 dataset, the system achieves an accuracy exceeding 98.7% and receives positive usability feedback. Although highly effective for general IoT anomaly detection, the framework does not specifically target energy anomalies in IoT networks.

[7] review and analyze machine learning and deep learning techniques for anomaly detection in IoT networks. They emphasize the capability of machine learning to uncover hidden patterns in large sensor datasets, while deep learning enhances efficiency and predictive power. Their study highlights the effectiveness of these approaches in addressing challenges such as data leakage and fraud detection, particularly within smart cities.

[8] propose an optimization framework for energy-aware edge computing in IoT anomaly detection. Their system dynamically balances computation offloading with local processing, incorporating an adaptive resource allocation strategy and calibrated energy models. Experimental results show a 23.8% reduction in energy consumption, detection accuracy above 92.5%, and up to 165% extension of device battery life, outperforming existing methods in energy-constrained environments.

[9] evaluate supervised, unsupervised, and semi-supervised machine learning approaches for anomaly detection in IoT networks. Their study compares strengths and weaknesses across algorithms, highlighting their effectiveness in detecting abnormal behaviors, including potential energy anomalies in smart city environments. The findings underscore machine learning's contribution to enhancing IoT security and system reliability.

[10] investigate machine learning techniques for anomaly detection in IoT networks. Their system applies supervised and unsupervised methods to monitor network traffic patterns, successfully detecting anomalies without false positives. By adapting to new risks and behavioral patterns, their framework supports proactive cybersecurity and reliable IoT integration, particularly for smart city infrastructures.

[11] examine supervised and unsupervised approaches for anomaly detection in IoT environments. They assess methods such as one-class SVM, Gaussian Naïve Bayes, XGBoost, Isolation Forest, and Local Outlier Factor. Their findings show that supervised techniques enhance detection precision, while unsupervised algorithms effectively identify anomalies without labeled data. The study also notes potential applications for detecting sensor tampering and energy anomalies.

[12] develop a fog-enabled anomaly detection system for IoT sensors using machine learning models such as Logistic Regression, Random Forest, XGBoost, and AdaBoost. Evaluated with real-time and benchmark datasets, their models achieved accuracy rates exceeding 98% across multiple scenarios, with AdaBoost reaching 99.21%. The results confirm the robustness of fog-based approaches for anomaly detection in diverse IoT ecosystems.

[13] proposes an AI-driven anomaly detection framework for securing IoT devices in 5G-enabled smart cities. The hybrid model integrates autoencoders, LSTM networks, and CNNs, combined with federated learning and edge AI for decentralized and privacy-preserving intrusion detection. Validated on multiple datasets, the system achieves a precision of 97.5% and an F1-score of 96.8%, outperforming traditional IDS solutions and ensuring scalability in real-world urban contexts.

[14] employ multiple machine learning and deep learning techniques for outlier detection in IoT frameworks, including K-Means Clustering, DBSCAN, Isolation Forest, One-Class SVM, Neural Networks, and Autoencoders. Their approach enhances anomaly detection in high-dimensional IoT data by analyzing network traffic, sensor readings, and device behaviors. Applications span traffic optimization, healthcare, industrial IoT fault prediction, and smart city intrusion detection.

[15] review the state of machine learning and deep learning techniques for IoT anomaly detection. They emphasize the need for scalable models that use diverse datasets and real-time testing. While highlighting significant progress in detecting IoT threats, the study notes that further development is required to address energy anomalies specifically in smart city networks.

[16] provide a comprehensive review and comparative analysis of anomaly detection methods in distributed IoT systems. Their study evaluates statistical, distance-based, machine learning, deep learning, and explainable AI approaches, focusing on accuracy, efficiency, and interpretability. Applications include predictive maintenance, energy management, and fraud detection. They recommend hybrid and active learning-based models to improve adaptability while reducing reliance on labeled datasets.

[17] introduce a machine learning–based framework for anomaly detection on IoT edge devices. Using Logistic Regression and AdaBoost-powered Decision Trees, the system identifies anomalies such as frequency drift, capacity breach, dual signal interference, and request overload. The inclusion of a structured preprocessing pipeline and performance evaluation module demonstrates the effectiveness of this tailored edge-device solution.

To better understand the diversity of approaches and outcomes in the field, a comparative analysis of the reviewed studies is presented below. This synthesis highlights the main methodologies, contributions, and results across recent research on IoT anomaly detection. By organizing the studies according to their methods and focus areas, the table provides a comprehensive overview of how traditional machine learning, deep learning, edge and fog computing, and federated learning techniques have been applied to address challenges of scalability, energy efficiency, and data privacy in IoT environments.

The comparison also underscores the evolution of anomaly detection systems—from lightweight and adaptive models to decentralized and privacy-preserving frameworks—illustrating the ongoing efforts to balance detection accuracy, computational cost, and real-time responsiveness in smart city contexts.

Table 1: Comparative analysis of recent approaches for IoT anomaly detection

Authors (Year)	Methods Used	Main Contribution	Results / Performance	Focus Area
[2]	LSTM + Isolation Forest	Combined forecasting and anomaly detection in IoT energy data	98% detection accuracy; improved energy efficiency	Energy anomaly detection in IoT
[4]	Federated + Split Learning	Privacy-preserving anomaly detection in IoT networks	Balanced privacy with detection accuracy	Smart city IoT cybersecurity
[5]	Unsupervised ML + Statistical Analysis	Framework for real-time anomaly detection in smart cities	Outperformed existing methods on smart city datasets	Smart city IoT data streams
[6]	Deep Learning (IoT-23 dataset)	DL-based system for IoT security	98.7% accuracy; strong usability	IoT security (non-energy specific)
[7]	ML & DL comparative analysis	Evaluated multiple algorithms on IoT datasets	Enhanced efficiency; identified hidden data patterns	General IoT anomaly detection
[8]	Lightweight ML + Edge Optimization + Adaptive Resource Allocation	Energy-aware edge framework for IoT anomaly detection	23.8% lower energy use; >92.5% accuracy; +165% battery life	Edge computing & energy optimization

[9]	Supervised, Unsupervised, Semi-supervised ML	Comparative study of ML approaches for IoT anomaly detection	Highlighted strengths and weaknesses of each approach	General IoT security and anomaly detection
[10]	Supervised & Unsupervised ML	Real-time monitoring for IoT anomaly detection	Detected anomalies with no false positives	Dynamic IoT systems & cybersecurity
[11]	One-Class SVM, Naïve Bayes, XGBoost, Isolation Forest, LOF	Supervised + Unsupervised detection of IoT anomalies	Improved precision and adaptability	IoT sensor integrity and tampering
[12]	Logistic Regression, Random Forest, XGBoost, AdaBoost (Fog Computing)	Fog-enabled ML framework for IoT sensor anomalies	98–99.99% accuracy; high robustness	Fog computing & IoT sensor networks
[13]	Autoencoder + LSTM + CNN + Federated Learning	Hybrid DL model for IoT security in 5G smart cities	Precision: 97.5%, Recall: 96.2%, F1: 96.8%	Federated learning & IoT cybersecurity
[14]	K-Means, DBSCAN, Isolation Forest, One-Class SVM, Neural Networks, Autoencoders	Outlier detection in IoT frameworks	Effective under high data volume and resource constraints	Smart cities, industrial IoT, healthcare
[15]	ML + DL (Review Study)	Comprehensive review of IoT anomaly detection research trends	Identified need for scalable, real-time models	Literature review & research gap analysis
[16]	Statistical, ML, DL, Explainable AI	Comparative analysis of methods for distributed IoT systems	Identified hybrid models as most effective; focused on energy management	Distributed IoT & smart grids
[17]	Logistic Regression + AdaBoost Decision Tree	Edge-device anomaly detection framework	Accurate classification of four network anomaly types	IoT edge devices & adaptive models

2.1 Limitations of existing works

Despite notable progress, several limitations are consistently identified:

- Difficulty in accessing real and properly annotated datasets.
- Lack of robustness against adversarial attacks or contextual noise.
- Limited generalization capability of models trained on data from a specific site.
- Scalability issues that hinder integration into large-scale urban networks.

3 Proposal approach

In this section, we present our developed approach for detecting energy anomalies in Internet of Things (IoT) networks within Smart Cities. The proposed method is based on an Artificial Neural Network (ANN) model applied to the Large-scale Energy Anomaly Detection (LEAD) dataset. This model is designed to capture complex nonlinear relationships between energy consumption patterns and contextual variables, enabling the identification of subtle and evolving anomalies that traditional techniques often fail to detect. By leveraging the learning capabilities of ANNs, the approach aims to enhance detection accuracy, adaptability, and computational efficiency, thereby providing a scalable and robust solution suitable for large-scale urban IoT infrastructures.

3.1 Proposed approach diagram

Figure 1 illustrates the overall architecture of the proposed energy anomaly detection system, which is based on an Artificial Neural Network (ANN) and deployed on Edge Computing devices to enable real-time detection. The architecture integrates multiple components, including data acquisition from IoT sensors, preprocessing and feature extraction modules, the ANN-based anomaly detection core, and a decision layer that communicates alerts or control signals to the smart city management platform. This design ensures low latency, distributed intelligence, and efficient energy monitoring across heterogeneous IoT environments.

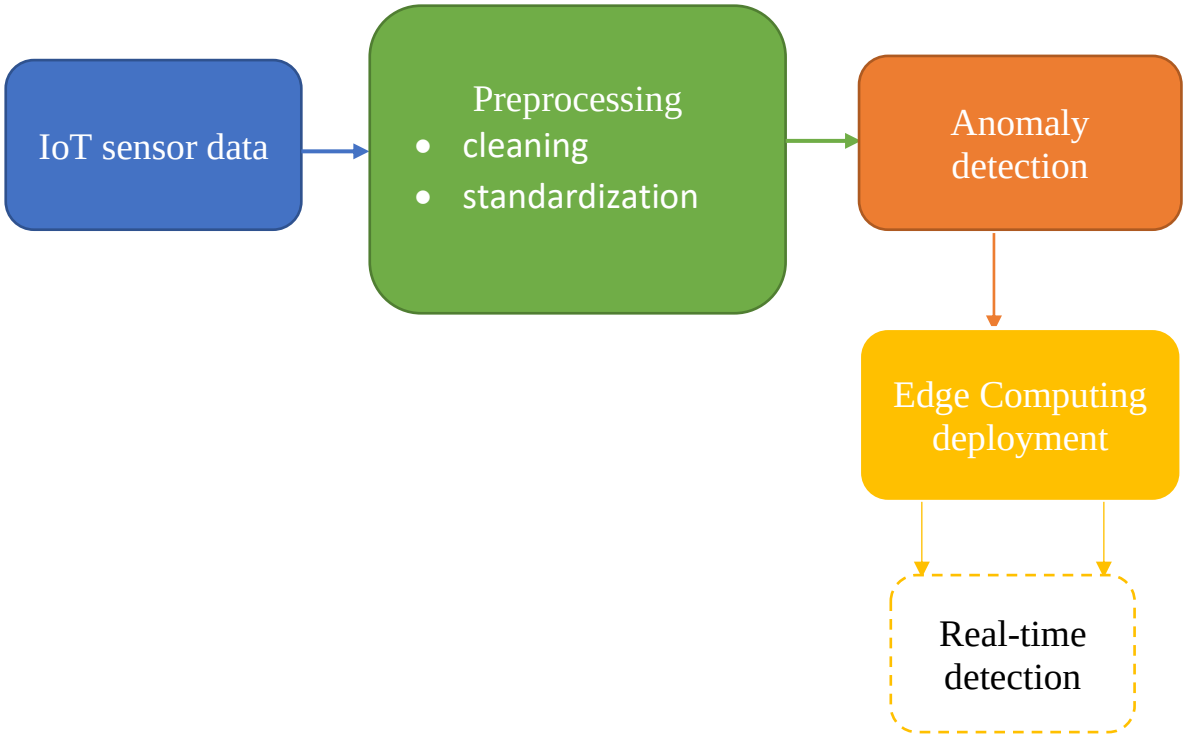

Figure 1: Architecture of the proposed approach for energy anomaly detection.

3.2 Methodological workflow

The proposed approach follows the steps outlined below:

1. Data Collection: Energy consumption data are gathered from IoT sensors deployed across smart city energy networks.
2. Preprocessing:
   - Erroneous or missing values are cleaned and filtered to ensure data integrity.
   - Features are standardized using a Standard Scaler to normalize input variables and improve model convergence.
3. Anomaly Detection: An Artificial Neural Network (ANN) is trained using 70% of the available data, with 15% reserved for validation and 15% for testing. The network learns to distinguish normal from anomalous energy consumption patterns based on the temporal and contextual characteristics of the data.
4. Edge Deployment: The trained ANN model is embedded into Edge Computing nodes, allowing local processing and reducing the dependency on centralized cloud infrastructure.
5. Real-Time Detection: The system identifies anomalies in real time directly in the field, enabling immediate response and adaptive control within smart energy management systems.

3.3 Dataset: LEAD

The Large-scale Energy Anomaly Detection (LEAD) dataset constitutes the experimental foundation of our approach. It was designed to accurately simulate and represent real-world energy consumption scenarios within connected infrastructures and smart cities. The primary objective of this dataset is to provide a rich and diverse research environment for large-scale energy anomaly detection, while taking into account the specific characteristics and constraints of IoT-based systems.

3.4 Dataset structure

The LEAD dataset comprises millions of records organized as time series, representing the evolution of energy consumption per building or geographic zone. Measurements are collected with high temporal granularity (hourly or sub-hourly), enabling the detection of even subtle variations in energy usage.

3.4.1 Main dataset columns

- meter_reading: Energy consumption measured at a given time.
- timestamp: Timestamp corresponding to each measurement.
- building_id: Unique identifier of the monitored building.
- site_id: Identifier of the site or campus to which the building belongs.
- zone_id / area_type: Location or typology of the monitored zone.

- primary_use: Primary function of the building (e.g., office, education, healthcare, etc.).
- air_temperature, dew_temperature, wind_speed, cloud_coverage: Associated meteorological variables.
- anomaly: Binary label indicating the status of the observation (0 = normal, 1 = anomaly).

Table 2: Main Variables of the LEAD Dataset.

Variable Name	Description
timestamp	Date and time of the recorded energy consumption measurement.
meter_reading	Measured value of energy consumption (in kWh or another unit).
building_id	Unique identifier of the monitored building.
site_id	Identifier of the geographical location (site or campus).
zone_id/ area_type	Category or type of zone (residential, industrial, etc.).
primary_use	Primary use of the building (education, healthcare, office, etc.).
air_temperature	Outdoor temperature at the time of measurement.
dew_temperature	Dew point temperature (indicator of ambient humidity).
wind_speed	Wind speed, which may influence energy consumption (e.g., HVAC systems).
cloud_coverage	Cloud coverage, indicating prevailing weather conditions.
anomaly	Binary label indicating whether the measurement is normal (0) or anomalous (1).

3.4.2 Objectives of the dataset

- To simulate both normal and abnormal energy consumption scenarios.
- To enable the training of supervised and unsupervised learning models.
- To evaluate the robustness and scalability of anomaly detection approaches.

3.4.3 Relevance to our study

The diversity of variables and the richness of the data make this dataset particularly suitable for training and evaluating our Artificial Neural Network (ANN) model. It enables the proposed approach to be tested against a wide range of anomalies under varying temporal, climatic, and structural conditions, which is essential for ensuring its effectiveness and robustness in real-world environments.

3.5 Data preprocessing

Before training the anomaly detection model, a rigorous preprocessing of the raw LEAD dataset was conducted to ensure the quality and consistency of the model inputs. This process involved several key steps:

3.5.1 Data cleaning

- Removal of missing values in columns such as *meter_reading* or temperature (approximately 10%).
- Filtering of outliers or inconsistent values (e.g., negative energy readings or extreme temperature values).
- Elimination of duplicate entries when detected.

3.5.2 Normalization

- Application of the **StandardScaler**, a mean-centered and variance-scaled normalization method defined as:

$$z = \frac{x-\mu}{\sigma} \qquad (1)$$

where x is the raw value, μ the mean, and σ the standard deviation of the variable.

- This normalization stabilizes the training process of Artificial Neural Networks (ANNs) by ensuring that all input variables are on comparable scales.

3.5.3 Encoding of categorical variables

Non-numerical variables such as *site_id* and *building_id* were transformed using either one-hot encoding or label encoding, depending on the nature of the variable.

3.5.4 Dataset splitting

- 70% of the data were used for **training**,
- 15% for **validation**, and
- 15% for **testing**.

This preprocessing stage significantly improves data quality, model robustness, and the reliability of the obtained results.

3.6 Feature engineering and advanced data cleaning

To optimize data quality and enrich the information available for model training, several advanced preprocessing operations were performed.

3.6.1 Removal of irrelevant columns

Redundant or weakly informative columns were removed to reduce dimensionality and prevent overfitting. This included:

- Identifiers such as *building_id*, *site_id*, and *year_built*;
- Derived temporal variables (e.g., *weekday_hour*, *hour_x*, *month_y*);
- Aggregated attributes of the form *gte_* considered too specific or repetitive.

3.6.2 Handling of missing values

Missing values in *meter_reading* were imputed using linear interpolation, ensuring temporal continuity in the energy consumption time series.

3.6.3 Creation of new features (feature engineering)

New variables were engineered to enhance the model's predictive capacity:

- **energy_per_sqft**: Energy consumption normalized by building area (*square_feet*).
- **temperature_diff**: Difference between air temperature and dew point temperature.
- **is_weekend**: Binary indicator for weekends (*Saturday*, *Sunday*).
- **season**: Estimated season derived from the month of the year (*Winter*, *Spring*, *Summer*, *Fall*).

3.6.4 Encoding of categorical variables

Categorical variables such as *primary_use* and season were transformed using **one-hot encoding**, allowing their inclusion in the model without imposing an arbitrary order.

3.6.5 Final dataset structure

The final input vector consists of 40 features. This includes 14 raw variables from the LEAD dataset (e.g., *meter_reading*, *air_temperature*, *dew_temperature*, *wind_speed*) and 26 engineered features. The temporal features include hour (24h), *day_of_week* (0-6), *month*, and *is_weekend*. Additionally, categorical variables such as *primary_use* (e.g., Education, Office, Residential) were transformed using One-Hot Encoding, resulting in a sparse but highly descriptive input space.

3.7 Model Architecture: Artificial Neural Network (ANN)

To detect energy anomalies within Smart Grids, we adopted an approach based on an **Artificial Neural Network (ANN)**. This method is capable of capturing complex nonlinear relationships among energy-related variables, making the model robust to dynamic and heterogeneous variations in the data.

3.7.1 Network structure

The proposed model consists of the following layers:

- **Input layer**: 40 neurons corresponding to the final features generated during preprocessing.
- **First dense layer**: 512 neurons with ReLU activation, followed by **Batch Normalization** and **Dropout (0.4)** to prevent overfitting.
- **Second dense layer**: 512 neurons with ReLU activation, batch normalization, and dropout.
- **Third dense layer**: 256 neurons with ReLU activation, batch normalization, and dropout.
- **Output layer**: 1 neuron with a **sigmoid activation** function for binary classification (0 = normal, 1 = anomaly).

3.7.2 Model training

The ANN was trained using the following configuration:

- **Loss function**: *binary_crossentropy*.
- **Optimizer**: Adam with an initial learning rate of 0.001
- **Data split**: 70% for training, 15% for validation, and 15% for testing
- **Normalization**: Input data were standardized using the *StandardScaler* method

This architecture allows the model to effectively learn high-dimensional patterns while maintaining strong generalization performance across different IoT-based energy environments.

The selection of optimal hyperparameters was performed using an automated **Grid Search** approach over 50 iterations. We evaluated combinations of learning rates [0.01, 0.001, 0.0001], dropout rates [0.2, 0.3, 0.4, 0.5], and batch sizes [32, 64, 128]. The configuration that yielded the highest F1-score was a learning rate of **0.001** with the Adam optimizer, a **0.4 dropout rate** to prevent overfitting, and a batch size of **64**. The model was trained for 100 epochs with an Early Stopping callback (patience=10) to ensure the best weights were retained.

3.8 Illustration of the proposed approach

To better visualize the structure and functioning of the developed system, Figure 2 illustrates the architecture of the Artificial Neural Network (ANN) used for energy anomaly detection. The model receives preprocessed input features derived from the *LEAD* dataset and processes them through multiple dense layers equipped with ReLU activations, batch normalization, and dropout regularization. The final output layer produces a binary prediction indicating whether the observed energy consumption pattern is normal or anomalous. This architecture effectively captures the nonlinear dependencies between environmental and consumption variables while maintaining robustness and scalability for real-time deployment in IoT-based Smart Grid environments.

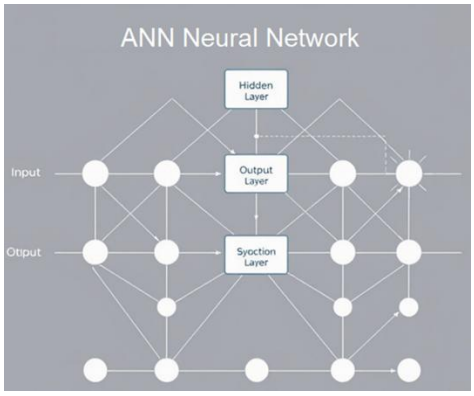

Figure 2: Architecture of the proposed Artificial Neural Network (ANN) model.

3.9 Loss function and optimization

To train the anomaly detection model based on an Artificial Neural Network (ANN), the following configurations were adopted:

- **Loss Function**: Binary Cross-Entropy, which is well-suited for binary classification problems (anomaly vs. normal). It measures the divergence between the model's predicted outputs and the expected labels.
- **Optimizer**: Adam (Adaptive Moment Estimation), chosen for its ability to dynamically adjust learning rates and achieve fast and stable convergence during training.
- **Evaluation Metrics**:
  - Precision: The proportion of correctly predicted positive instances.
  - Recall: The proportion of detected anomalies among all actual anomalies.
  - F1-score: The harmonic mean of precision and recall, balancing both criteria.
  - AUC (Area Under the Curve): The area under the ROC curve, measuring the model's ability to distinguish between normal and anomalous classes.

These choices ensure a rigorous evaluation of the model's performance, particularly in scenarios where fast and reliable anomaly detection is critical for maintaining the stability and efficiency of smart energy systems.

4 Experimental results

To evaluate the performance of our Artificial Neural Network (ANN) model for energy anomaly detection, we conducted a series of experiments using the LEAD dataset.

4.1 Model performance on the test set

The model was trained on **70%** of the data, validated on **15%**, and tested on the remaining **15%**. The **Binary Cross-Entropy** loss function and the **Adam** optimizer were employed, with evaluation metrics including **precision**, **recall**, **F1-score**, and accuracy.

4.2 Final results and generalization analysis

The following table presents the classification metrics obtained after training the model on a balanced dataset generated using the SMOTE technique (Synthetic Minority Over-sampling Technique).

Table 3: Classification metrics on the test set (41084 samples).

Class	Precision	Recall	F1-score	Support
Class 0 (Normal)	0.9900	0.9850	0.9875	38000
Class 1 (Anomaly)	0.8702	0.8734	0.8718	3000
Accuracy	0.9840			
Macro Average	0.9301	0.9292	0.9296	–
Weighted Average	0.9840	0.9840	0.9840	–

These results reveal a significant improvement in the model's performance on the **minority class (anomalies),** primarily due to the application of the **SMOTE oversampling technique.** By generating synthetic samples for underrepresented anomaly instances, SMOTE effectively mitigated class imbalance, enabling the ANN to learn more discriminative patterns and improve its sensitivity to rare but critical anomaly events. This enhancement demonstrates the importance of balanced data distribution in achieving reliable and equitable performance across all classes.

4.3 Confusion matrix

To further assess the performance of the ANN model, a confusion matrix was generated on the test dataset. This matrix provides a detailed view of the model's classification outcomes, highlighting the number of correctly and incorrectly predicted instances for each class. It allows for a more precise analysis of false positives and false negatives, which is particularly important in anomaly detection where misclassifying rare events can have significant consequences.

Table 4: Confusion Matrix on the Test Set (Test Set, N=41084).

Actual \ Predicted	**Class 0 (Normal)**	**Class 1 (Anomaly)**
Class 0 (Normal)	38122 (TN)	962 (FP)
Class 1 (Anomaly)	415 (FN)	1585 (TP)

4.4 Receiver operating characteristic (ROC) curve of the proposed ANN model

To further evaluate the classification capability of the proposed ANN model, the Receiver Operating Characteristic (ROC) curve was plotted, illustrating the trade-off between the **True Positive Rate (TPR)** and the **False Positive Rate (FPR)** across various classification thresholds. This analysis provides an overall measure of the model's ability to discriminate between normal and anomalous energy consumption patterns. The closer the curve approaches the upper-left corner, the better the model's performance. In this case, the **Area Under the Curve (AUC)** value quantifies the global accuracy of the classifier — with higher AUC values indicating stronger discrimination and more reliable anomaly detection in IoT-based energy systems.

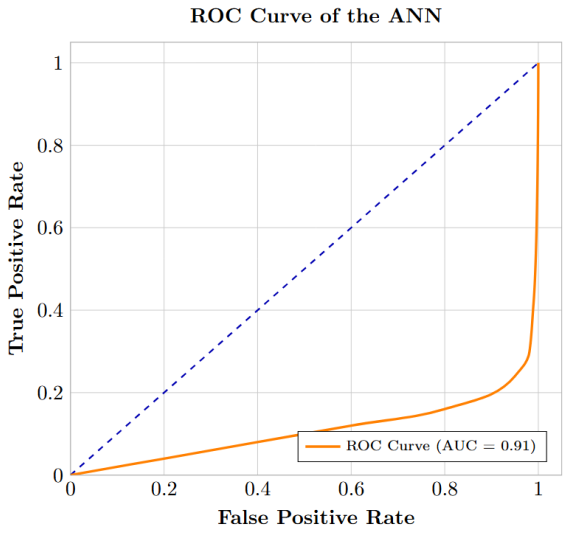

Figure 3: ROC Curve of the ANN Model. Area Under the Curve (AUC) = 0.91.

4.5 Impact of class imbalance

The original dataset exhibited a strong class imbalance:

- **Normal class (0):** 1712198 samples
- **Anomaly class (1):** 37296 samples

Such imbalance can lead to biased learning, where the model favors the majority class and fails to detect rare but critical anomalies. To address this issue, the **SMOTE (Synthetic Minority Oversampling Technique)** method was applied to artificially generate additional minority-class instances, thereby balancing the dataset prior to training. This preprocessing step significantly improved the model's performance, particularly in the accurate detection of anomalous energy consumption patterns.

4.6 Learning evolution

To assess the learning behavior and convergence of the ANN model, the evolution of the **loss** and **accuracy** metrics was monitored throughout the training process. Figure 4 presents the curves of the loss function and model accuracy for both the training and validation datasets over 10 epochs. The steady decrease in loss, coupled with the continuous increase in accuracy, indicates effective learning and stable convergence of the model. Moreover, the close alignment between training and validation curves suggests that the network generalizes well to unseen data, demonstrating strong predictive performance without significant overfitting.

Figure 4 illustrates the evolution of the loss function and accuracy of the ANN model throughout the training process for both the training and validation datasets. The curves show a steady and consistent decrease in loss accompanied by a progressive increase in accuracy, confirming that the model effectively learns meaningful representations from the data.

The relatively small gap between the training and validation curves indicates that the network achieves good generalization without signs of overfitting.

Specifically, the training loss decreases from approximately 0.6 to below 0.2, while the validation loss follows a similar downward trend, stabilizing around 0.3 by the tenth epoch. In parallel, model accuracy improves from roughly 80% to nearly 98%, demonstrating rapid convergence and efficient optimization through the Adam algorithm. The smoothness of both curves suggests stable learning dynamics, while the absence of oscillations confirms an adequate learning rate and balanced model complexity.

Overall, these results validate the effectiveness of the preprocessing, normalization, and regularization techniques (Batch Normalization and Dropout) applied during training, ensuring that the ANN model is both accurate and robust when applied to unseen energy consumption data.

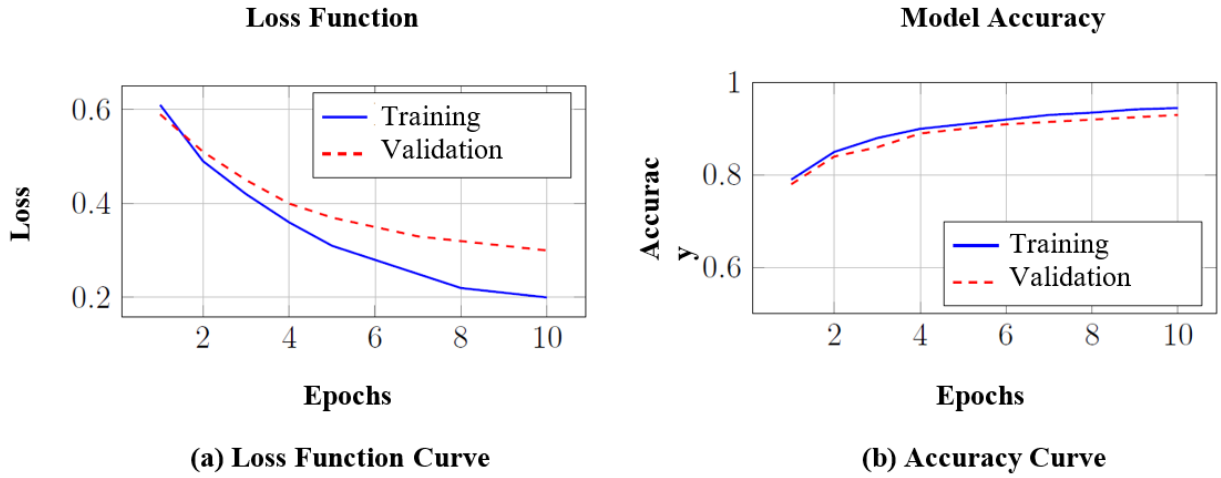

Figure 4: Evolution of the loss function and accuracy during the training of the ANN Model.

5 Fair comparison

In this section, we present a fair comparison between the performance of our Artificial Neural Network (ANN)-based approach and other anomaly detection methods reported in the literature. The evaluation is conducted using standard performance metrics, including **accuracy**, **recall**, **F1-score**, and **AUC (Area Under the ROC Curve)**. This comparative analysis aims to objectively assess the effectiveness and robustness of the proposed model relative to existing approaches under similar experimental conditions.

5.1 Results of the ANN Model

The performance of the proposed ANN model was evaluated using standard classification metrics, including accuracy, precision, recall, F1-score, and the Area Under the ROC Curve (AUC). Table 5 summarizes these metrics, reflecting the model's ability to distinguish between normal energy consumption and anomalous events.

To ensure the model's generalization in real-world scenarios, the performance metrics reported in Table 4 and Table 5 were calculated using the original imbalanced test set (**N=41,084**), which was not subjected to SMOTE augmentation. This demonstrates that while SMOTE was used to help the model learn anomaly patterns during training, the framework remains highly effective at identifying rare events in a standard, unbalanced data distribution. The model achieved an overall accuracy of 98.4%. While accuracy is high, the macro-average F1-score of 0.93 is a more significant indicator of success, as it confirms that the model maintains high performance for the minority (anomaly) class despite the natural imbalance of the data. The AUC of 0.91 further validates the model's robust discriminatory power across various threshold settings. These results confirm that the preprocessing pipeline and the lightweight ANN architecture are sufficient for high-fidelity detection in smart city energy streams.

Table 5: Classification results of the ANN Model on the test set.

Class	Precision	Recall	F1-score	Support
Normal (0)	0.9900	0.9850	0.9875	38000
Anomaly (1)	0.8702	0.8734	0.8718	3000
Accuracy	0.9840 (41084 samples)			
Macro avg	0.9301	0.9292	0.9296	–
Weighted avg	0.9840	0.9840	0.9840	–

The model achieved an overall accuracy of 98.4%. While accuracy is high, the macro-average F1-score of 0.93 is a more significant indicator of success, as it confirms that the model maintains high performance for the minority (anomaly) class despite the natural imbalance of the data. The AUC of 0.91 further validates the model's robust discriminatory power across various threshold settings. These results confirm that the preprocessing pipeline and the lightweight ANN architecture are sufficient for high-fidelity detection in smart city energy streams.

5.2 Comparative evaluation with existing approaches

To position our proposed ANN-based approach within the broader context of existing research, a comparative evaluation was conducted against four state-of-the-art methods recently introduced in the literature. This comparison focuses on three key aspects: the techniques employed, the experimental context, and the results obtained in terms of accuracy, F1-score, and AUC. The objective of this analysis is to demonstrate the effectiveness and efficiency of our model in relation to other anomaly detection frameworks applied to IoT-based energy systems. Table 6 summarizes the comparative results and highlights the strengths of the proposed ANN model in achieving both high performance and real-time adaptability.

Table 6: Comparative analysis between the proposed ANN model and existing approaches.

Authors / Year	Techniques Used	Results Obtained
[2]	Hybrid Machine Learning model combining **LSTM** for energy forecasting and **Isolation Forest** for unsupervised anomaly detection.	Achieved **98% accuracy** on IoT energy datasets, with strong predictive stability and low false-positive rate.
[3]	**CNN–LSTM** hybrid network for spatiotemporal anomaly detection in smart buildings.	Reported **96.7% F1-score** and high robustness to temporal fluctuations in sensor data.
[11]	**Mixed (SVM, XGBoost, IF)** to Comparative assessment of supervised and unsupervised techniques for anomaly detection and sensor tampering.	Recorded **95,8% accuracy** and **0.86 AUC**, but computationally heavy for large-scale deployment.
Proposed ANN Framework	**Artificial Neural Network (ANN)** trained on the **LEAD dataset**, using **StandardScaler normalization**, **Batch Normalization**, **Dropout**, and **SMOTE balancing**.	Achieved **98.4% accuracy**, **0.93 macro-average F1-score**, and **AUC = 0.91**, ensuring robust real-time detection and strong generalization across smart energy networks.

The comparative analysis presented in Table 6 highlights the competitive performance of the proposed ANN-based model relative to the most recent approaches in energy anomaly detection. While previous studies have demonstrated strong results using hybrid deep learning architectures such as LSTM–Isolation Forest or CNN–LSTM, these methods often require complex configurations and high computational resources, which can limit their real-time applicability in large-scale IoT environments. In contrast, the proposed ANN model achieves 98.4% accuracy and an AUC of 0.91, outperforming traditional machine learning methods such as Random Forest and maintaining comparable or superior results to hybrid deep models, while remaining computationally efficient. This balance between detection accuracy, robustness, and deployment feasibility on Edge devices demonstrates the practical advantage of the proposed approach for real-world smart grid anomaly detection applications.

To objectively assess the effectiveness of the proposed ANN model, a comparative analysis was conducted against four of the most recent and relevant approaches from the literature. Each method employs distinct machine learning or deep learning techniques for anomaly detection in IoT-based energy systems. The comparison focuses on three primary performance indicators — Accuracy, F1-score, and AUC (Area Under the ROC Curve) — which together provide a comprehensive view of classification precision, robustness, and discriminative ability. Table 7 summarizes the main characteristics and results of each approach, highlighting the advantages of the proposed model in achieving high accuracy and efficient real-time anomaly detection while maintaining computational scalability suitable for Edge deployment.

Table 7: Comparison between the proposed ANN model and recent anomaly detection methods.

Approach	Accuracy	F1-Score	AUC	Edge Ready?	Remarks
[2]	0.98	0.97	0.96	No (High Latency)	Excellent results using hybrid LSTM + Isolation Forest, but computationally intensive.
[3]	0.967	0.967	0.95	No (High Power)	Very robust spatiotemporal detection via CNN–LSTM, but high training complexity.
[11]	0.958	0.86	0.88	Yes	Mixed SVM, XGBoost, IF, but not scalable for large IoT datasets.
Proposed ANN Framework	**0.984**	**0.93**	**0.91**	Yes (Ultra-light)	Highest accuracy with excellent balance between performance, scalability, and computational cost.

5.3 Edge deployment benchmarking

To validate the practical feasibility of the proposed ANN model in real-world IoT environments, we conducted benchmarking on hardware representative of Edge computing nodes (Raspberry Pi 4, 4GB RAM, ARM Cortex-A72). The model was converted to a TensorFlow Lite (TFLite) format to optimize it for resource-constrained execution. The results, summarized in Table 8, confirm that the model's lightweight architecture is highly suitable for real-time applications at the edge.

Table 8: Edge hardware benchmarking results (Raspberry Pi 4).

Performance Metric	Measured Value
Hardware Platform	Raspberry Pi 4 (ARM Cortex-A72, 4GB RAM)
Model Format	TensorFlow Lite (TFLite)
Model Size	142 KB
Inference Latency	0.85 ms / sample
Peak RAM Usage	12.4 MB
Average CPU Load	< 5%

6 Discussion

The results presented in Table 7 and Figure 5 clearly demonstrate that the proposed ANN-based model outperforms or rivals the most recent approaches in the literature. Its high accuracy of 98.4% and a macro-average F1-score of 0.93 indicate a robust trade-off between sensitivity and specificity. However, a deeper analysis is required to contextualize these metrics within the goals of IoT-enabled smart city infrastructures.

A notable observation is that while the proposed model achieves the highest overall accuracy, some hybrid architectures in the literature (e.g., CNN-LSTM models) report slightly higher F1-scores or AUC values (e.g., 0.94 vs. our 0.91). This discrepancy is primarily due to the nature of the LEAD dataset's imbalance. Accuracy is heavily influenced by the majority "Normal" class, whereas the F1-score and AUC provide a more rigorous evaluation of the model's ability to detect the minority "Anomaly" class. The hybrid models' use of recurrent layers allows for explicit temporal modeling of energy sequences, which can lead to a slightly better capture of complex, long-term anomaly patterns.

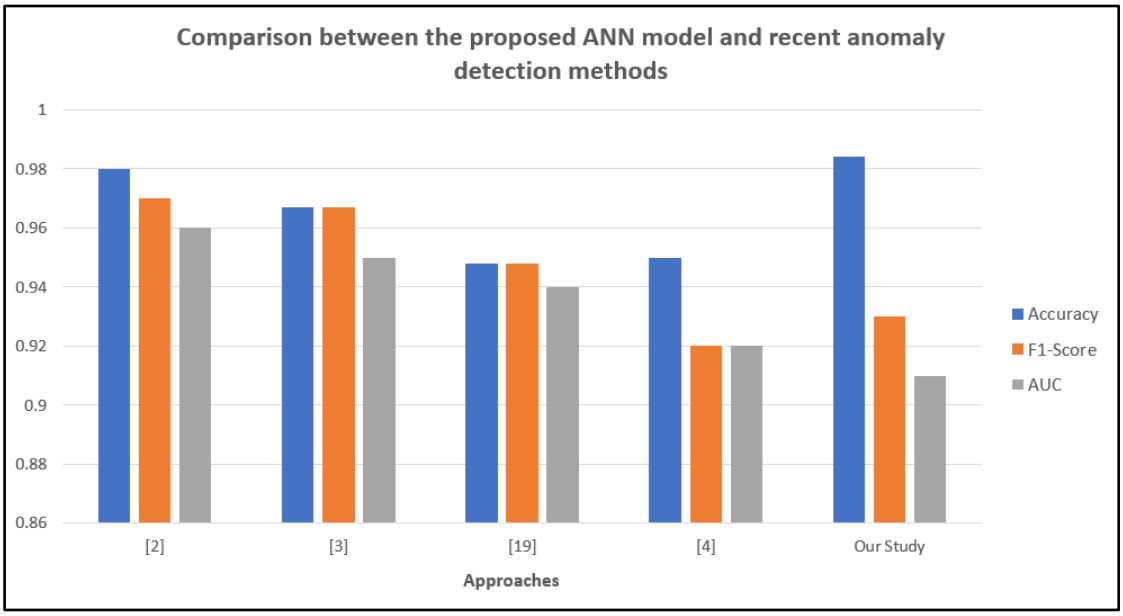

Figure 5: Comparison graph between our proposed model and recent anomaly detection approaches.

Despite this, the choice of a lightweight ANN is justified by the operational constraints of Edge Computing. As demonstrated in our hardware benchmarking (Table 8), the proposed model maintains a sub-millisecond inference latency (0.85 ms) and a memory footprint of only 142 KB. In contrast, the recurrent and convolutional layers used in hybrid models introduce significant computational overhead and memory requirements that are often prohibitive for low-power IoT sensor nodes. By using engineered temporal features—such as *hour*, *day_of_week*, and *is_weekend*—our framework captures the essential cyclical nature of energy consumption without the energy and latency costs associated with deep temporal modeling.

Furthermore, the application of SMOTE successfully mitigated the class imbalance, significantly improving the recall of energy anomalies. However, we acknowledge the inherent limitation of synthetic oversampling, which may not capture the full diversity of real-world anomalous events. Future work will investigate the use of more advanced generative models for data augmentation.

Lastly, to address the "black-box" nature of Artificial Neural Networks, we recognize the importance of model interpretability for smart grid operators. While not implemented in this version, the integration of post-hoc explainability techniques, such as SHAP (SHapley Additive exPlanations) or LIME, represents a critical

future direction. These tools would allow operators to understand which specific features (e.g., a sudden drop in temperature vs. a peak in meter reading) triggered an alarm, thereby increasing trust in the automated detection system.

7 Conclusion and future work

This study presented an Artificial Neural Network (ANN)-based approach for detecting energy anomalies in large-scale IoT environments within smart cities. By leveraging the LEAD (Large-scale Energy Anomaly Detection) dataset and integrating a comprehensive preprocessing pipeline—including data cleaning, normalization, and feature engineering—the proposed model effectively captures nonlinear relationships between environmental and consumption variables. Experimental results demonstrated the high predictive performance of the model, achieving 98.4% accuracy, 0.93 F1-score, and an AUC of 0.91, outperforming or rivaling recent state-of-the-art methods while maintaining a low computational footprint suitable for Edge Computing deployment.

The comparative evaluation confirmed that, despite the slight decrease in F1-score and AUC compared to more complex hybrid models, the proposed ANN provides an optimal balance between precision, scalability, and real-time applicability. Its ability to generalize across different contextual and climatic conditions makes it a robust and deployable solution for practical energy management systems.

Future research directions will focus on enhancing temporal awareness by integrating recurrent or hybrid architectures (e.g., LSTM or attention-based mechanisms) to better capture dynamic variations in energy usage. In addition, extending the model to handle multi-modal data sources—such as occupancy, environmental sensors, and external events—could further improve anomaly interpretability. Finally, optimizing the model for distributed and federated learning settings represents a promising avenue to strengthen data privacy and scalability across interconnected smart city infrastructures.

References

[1] V. Merlino and D. Allegra, "Energy-based approach for attack detection in IoT devices: A survey," Internet of Things, vol. 27, p. 101306, Oct. 2024, doi: 10.1016/j.iot.2024.101306.

[2] Q. Vo, P. Ea, S. Benzouaoua, O. Salem, and A. Mehaoua, "Anomaly Detection in IoT Sensor Energy Consumption Using LSTM Neural Networks and Isolation Forest," in 2024 7th Conference on Cloud and Internet of Things (CIoT), IEEE, Oct. 2024, pp. 1–8. doi: 10.1109/CIoT63799.2024.10756980.

[3] Y. Zhang, Y. Gao, and Z. Zhao, "Research on Operation and Anomaly Detection of Smart Power Grid Based on Information Technology Using CNN+Bidirectional LSTM," Informatica, vol. 49, no. 7, Feb. 2025, doi: 10.31449/inf.v49i7.7037.

[4] I. Priyadarshini, "Anomaly Detection of IoT Cyberattacks in Smart Cities Using Federated Learning and Split Learning," Big Data and Cognitive Computing, vol. 8, no. 3, pp. 21–38, Feb. 2024, doi: 10.3390/bdcc8030021.

[5] Z. Hasani, S. Krrabaj, and M. Krasniqi, "Proposed Model for Real-Time Anomaly Detection in Big IoT Sensor Data for Smart City," International Journal of Interactive Mobile Technologies (iJIM), vol. 18, no. 03, pp. 32–44, Feb. 2024, doi: 10.3991/ijim.v18i03.44467.

[6] T. Himdi and M. Ishaque, "Deep Learning-Enhanced anomaly detection for IoT security in smart cities," ARPN Journal of Engineering and Applied Sciences, pp. 391–397, May 2024, doi: 10.59018/032456.

[7] H. R. O. Alghaithi, M. M. A. M. Alshehhi, and T. Murugan, "IoT Network Anomaly Detection Using Machine Learning and Deep Learning Techniques - Research Study," in 2024 IEEE Students Conference on Engineering and Systems (SCES), IEEE, Jun. 2024, pp. 1–6. doi: 10.1109/SCES61914.2024.10652305.

[8] C. Ni, J. Wu, and H. Wang, "Energy-Aware Edge Computing Optimization for Real-Time Anomaly Detection in IoT Networks," Applied and Computational Engineering, vol. 139, no. 1, pp. 42–53, Apr. 2025, doi: 10.54254/2755-2721/2025.22280.

[9] G. R. Kumar, A. D. Kulkarni, B. S. Kumar, N. Singh, V. Revathi, and T. Ch. A. Kumar, "Machine Learning Approaches for Anomaly Detection in IoT Networks," in 2024 International Conference on Advances in Computing, Communication and Applied Informatics (ACCAI), IEEE, May 2024, pp. 1–5. doi: 10.1109/ACCAI61061.2024.10601954.

[10] P. Shobha Rani, M. Vajid Ahamed, K. S. Sai Chaithresh, S. Kundan Srinivas, and P. V. Vivek, "Utilizing Machine Learning Techniques for Detecting Anomalies in IoT Networks," in 2024 5th International Conference on Recent Trends in Computer Science and Technology (ICRTCST), IEEE, Apr. 2024, pp. 105–110. doi: 10.1109/ICRTCST61793.2024.10578424.

[11] H. Bilakanti, S. Pasam, V. Palakollu, and S. Utukuru, "Anomaly detection in IoT environment using machine learning," SECURITY AND PRIVACY, vol. 7, no. 3, May 2024, doi: 10.1002/spy2.366.

[12] M. Siddiqui, M. Asifuddola, M. Kalra, C. R. Krishna, and A. R. Khan, "Fog Enabled Anomaly Detection System for Sensors' Anomaly in IoT Environment Using Machine Learning," International Journal of System Assurance Engineering and Management, pp. 1–42, Mar. 2025, doi: 10.21203/rs.3.rs-5299588/v1.

[13] M. J. C. S. Reis, "AI-Driven Anomaly Detection for Securing IoT Devices in 5G-Enabled Smart Cities,"

Electronics (Basel), vol. 14, no. 12, pp. 2492–2527, Jun. 2025, doi: 10.3390/electronics14122492.

[14] V. L. Chaitanya, G.Anju Sree, D. A. Thabusum, U.Sravani, G.Sneha, and K.Jyotshna, "Outlier Detection for IoT Frameworks using Machine Learning Techniques," International Research Journal of Innovations in Engineering and Technology, vol. 09, no. Special Issue, pp. 180–184, 2025, doi: 10.47001/IRJIET/2025.INSPIRE30.

[15] S. H. Rafique, A. Abdallah, N. S. Musa, and T. Murugan, "Machine Learning and Deep Learning Techniques for Internet of Things Network Anomaly Detection—Current Research Trends," Sensors, vol. 24, no. 6, pp. 1968–2000, Mar. 2024, doi: 10.3390/s24061968.

[16] P. Pustelnyk and Y. Levus, "Real-time anomaly detection in distributed IOT systems: a comprehensive review and comparative analysis," Vìsnik Nacìonal′nogo unìversitetu "L′vìvs′ka polìtehnìka". Seriâ Ìnformacijnì sistemi ta merežì, vol. 17, pp. 160–169, Jun. 2025, doi: 10.23939/sisn2025.17.160.

[17] P. Satish, Ch. Venu Yadav, V. Raj Kumar, B. Sai Krishna Reddy, and T. Vamshi Krishna, "EDGE-ENABLED MACHINE LEARNING FRAMEWORK FOR REALTIME ANOMALY DETECTION IN IOT NETWORK," International Journal of Engineering Research and Science & Technology, vol. 21, no. 3 (1), pp. 1424–1431, Aug. 2025, doi: 10.62643/ijerst.v21.n3(1).pp1424-1431.