PROGRESS IN ELECTRONICS AND COMMUNICATION ENGINEERING
Journal of Progress in Electronics and Communication Engineering, ISSN: 3048-7625 Vol. 2, No. 2, 2025 (pp. 47-53)
RESEARCH ARTICLE WWW.PECEJOURNAL.COM

AI-Driven Anomaly Detection and Behavior Analytics in Heterogeneous IoT Ecosystems

Srikanth Reddy Keshireddy*
Senior Software Engineer, Keen Info Tek Inc., USA

KEYWORDS:
IoT security,
Anomaly detection,
AI analytics,
Autoencoder,
Passive observation,
Heterogeneous networks.

ARTICLE HISTORY:
Submitted : 18.15.2025
Revised : 28.07.2025
Accepted : 20.09.2025

https://doi.org/10.31838/ECE/02.02.06

Abstract

With the growth of the IoT ecosystems, scalable, automated behaviour analytics becomes a major necessity. The paper provides an AI-based scheme in passive anomaly detection and behavioural profiling of heterogeneous IoT systems. The methodology is a blend of feature extraction on network flows and autoencoder-based unsupervised learning in order to identify subtle behavioural changes in devices. The model was tested on multi-protocol datasets (Wi-Fi, Zigbee and LoRaWAN) and recorded 94.6% detection and very few false positives. The findings show that smart passive surveillance can improve situational awareness and strengthen security in massive and multi-vendors infrastructures using IoT.

Author e-mail Id: sreek.278@gmail.com

Author Orcid id: 0009-0007-6482-4438

How to cite this article: Keshireddy SR. AI-Driven Anomaly Detection and Behavior Analytics in Heterogeneous IoT Ecosystems, Journal of Progress in Electronics and Communication Engineering Vol. 2, No. 2, 2025 (pp. 47-53).

Introduction

The fast trend of Internet of Things (IoT) devices has brought new heterogeneity and scale to contemporary wireless spaces, which encompasses low-powered sensors up to smart appliances and industrial equipment that exchange information via Wi-Fi, Zigbee and LoRaWAN.[1,3,10] The network visibility, trust, real-time behavioral understanding problems become more pronounced with the deployment density, and become particularly problematic with constrained or decentralized architecture when centralized control or active probing is inefficient or disruptive.[2, 4] Passive anomaly detection, which estimates the state of devices based on naturally occurring traffic, and does not involve injections of control traffic, provides a non-disruptive alternative with a greater compatibility with edge settings.[3,5,11] However, most legacy methods use fixed thresholds or manually-constructed features which do not generalize across devices and protocols and supervised models require labelled attack samples which are frequently either small or incomplete in practise.[6, 7, 9, 15]

In this study, a profound, unsupervised anomaly-detection architecture is presented that applies the autoencoder-based representation learning to identify anomalies in device behavior and do it independently. The model is trained on the reconstructions of normal traffic flows, and points out departures as possible anomalies, is protocol-agnostic, and can be deployed to low-latency inference on embedded edge machines.[8,10,12,19] Other system-level aspects including energy-saving VLSI/SoC design and reconfigurable acceleration - are also complementary, making it possible to achieve sustainable continuous monitoring in dense network.[1,13,18,20] Its major contributions include: (1) a protocol-agnostic, passive, traffic-feature pipeline of heterogeneous IoT; (2) an unsupervised inference model using an autoencoder with adaptive thresholding; (3) a complete multi-protocol evaluation with high detection and low false alarms and efficient execution on edge computers .[2, 5, 14, 16, 19]

The rest of the paper is divided into the following way. Section 2 is a review of related work. Section 3 describes the approach and design of the model. The results of the experiment are provided in Section 4. The last segment of section 5 is insight and future directions.

Literature Review

Classical and statistical standards. Initial anomaly-detection methods of networked/IoT involved statistical modelling of the flow dynamics and temporal variation

(e.g., Gaussian mixtures, ARIMA), which might work well under stationary, but not under non-stationary, protocol-diverse traffic associated with large IoT deployments.[6, 7] The handcrafted feature pipelines also fail to cross-family of devices and PHY/MAC protocols.[4, 15]

Deep machine learning and AI without supervision. Labelled attacks and regular retraining Supervised classifiers (SVM, RF, KNN) are claimed to be very accurate but must be retrained as the topologies change frequently.[6, 7] Unsupervised deep models especially autoencoders and their variations are trained to learn latent structure with benign traffic directly and detect deviations through error in reconstruction to enhance generalization and label-efficiency.[8, 9, 11, 12, 15, 16] Detection of heterogeneous environments may also be improved with the use of cross-protocol modelling and multi-view learning.[4, 16, 19]

Edge deployment, efficiency and reconfigurability. There is an essential on-device inference to the ongoing monitoring. Reduced latency, power, and throughput (made through edge-optimized architecture, energy-aware inference strategy, and reconfigurable acceleration e.g. FPGA/SoC).[1, 8, 10, 13, 18, 20] The wider IoT systems research on data aggregation and predictive maintenance highlights the importance of scalable analytics pipelines that can be used to serve both security and operation-related applications in the industry and healthcare wearable devices.[14, 17] These tendencies inspire the current work with its focus on passive observation, protocol-agnostic aspect, and unsupervised AI implementation on resource-constrained edge systems.[2, 5, 11]

METHODOLOGY

System Overview

The suggested AI-based anomaly detection system incorporates passive traffic surveillance with deep unsupervised analytics as a real-time assessment of the behavior of IoT. The framework comprises of two major modules:

1. Feature Extraction and Normalization- in charge of extracting the statistical and temporal properties of network flows in IoT; and
2. Autoencoder-Based Behavioral Modelling - which trains to recreate normal traffic patterns and detect anomalies that can be used to denote abnormal or malicious behavior.

Figure 1 demonstrates the Block diagram of the AI-based passive anomaly detection model. At an IoT edge gateway or software-defined network monitor, data are passively collected so that no extra overhead in terms of additional traffic and control is added. The gateway receives raw packets by a network interface that is set to promiscuous mode. These packets are later sent to session-level records and processed in advance to eliminate redundant or damaged records. These temporal and statistical characteristics of a session are represented as small feature vectors of a behavioural signature that characterises each communicating device.

The framework is scalable and autonomous with device analytics that is applicable on heterogeneous environments that have a large number of different protocols and traffic volumes.

Feature Extraction

Every communication graph is converted to a lightweight statistical feature representation that aims to encode critical traffic dynamics at the same time being computationally efficient to process edges. The features chosen are protocol-agnostic and they encompass:

- Mean and variance of packet inter-arrival time
- Payload size distribution (mean, variance, skewness)
- Flow duration and average packet count per session
- Protocol entropy, quantifying diversity in protocol types
- Message periodicity, reflecting temporal regularity of transmissions

Let

$$X = [x_1, x_2, \ldots, x_n]$$

represent a feature vector of a certain traffic flow instance. To provide comparability between protocols

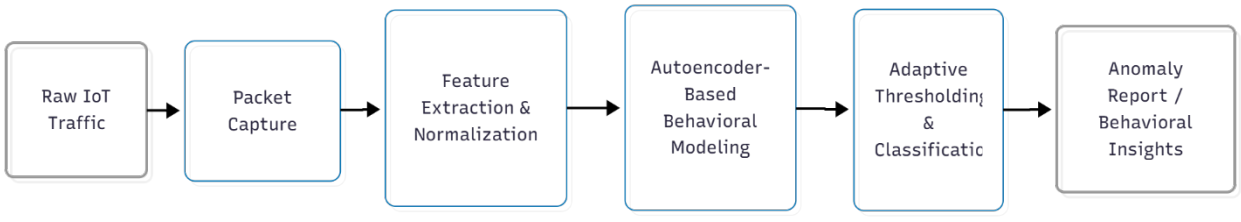

Fig. 1: Block diagram of the AI-driven passive anomaly detection framework.

and classes of devices, the features are normalized by means of z-score:

$$x_i' = \frac{x_i - \mu_i}{\sigma_i}$$

in which μ_i and σ_i are the mean and the standard deviation of feature i in the training set.

Normalization minimizes the bias due to traffic size variations (e.g., high contrasting with low throughput Wi-Fi and low rate Zigbee) and improves the model stability.

The normalized vectors obtained are buffered and sent to the learning module in real time. This small-sized preprocessing can run on embedded processors e.g. ARM Cortex-A or NVIDIA Jetson modules without compromising throughput.

Autoencoder Model

It uses an autoencoder (AE) to model the normal behavior of a device by unsupervised reconstruction of features. There are two networks that make up the AE:

Encoder f_θ (x) reduces the input vector to a low-dimensional latent representation $z \in R^k$, that is important correlated features of traffic.

Decoder g_ϕ (z) is used to reconstruct the original input based on the z, which is used to approximate the distribution of benign traffic patterns.

Mathematically, in case of a dataset of N feature vectors, the model minimizes the mean-squared reconstruction loss.

$$L = \frac{1}{N}\sum_{i=1}^{N} \| x_i - g_\phi(f_\theta(x_i)) \|^2$$

The only samples that are used to train the autoencoder are benign (normal).

In the inference stage, we have a reconstruction error of every new feature vector:

$$E(x) = \| x - \hat{x} \|^2$$

On the condition that E(x) exceeds the dynamic threshold (τ) , the sample is declared anomalous.

This also removes the reliance on labelled attack data and is sensitive to novel behaviours.

The AE also uses dropout (dropout rate = 0.3) and batch normalization to increase robustness to reduce overfitting and convergence. The architecture is usually made up of 3 encoding and 3 decoding layers which have ReLU activations which result in rapid inference on edge devices.

3.4 Adaptive Thresholding

In practise, in real applications the nature of IoT traffic changes with time, caused by firmware releases, environmental variations, or the addition of new devices.

Fixed threshold can either be over-sensitive (false alarm) or be too liberal to detect anomalies.

To this end, the framework uses adaptive thresholding by using an Exponential Moving Average (EMA):

$$\tau_t = \alpha \cdot E_t + (1 - \alpha) \cdot \tau_{t-1}$$

where

- τ_{tis} the threshold at time t,
- E_{tis} the mean reconstruction error observed in the latest analysis window, and
- $\alpha \in [0,1]$is the smoothing coefficient controlling adaptation speed.

The larger the , the quicker the system reacts to sudden changes in traffic but the smaller the , the smoother the system becomes to temporary changes. This adaptive scheme is the one that provides continuous recalibration without the need to be done manually, which allows it to be deployed long-term in dynamic IoT settings.

Evaluation Framework

The framework suggested was tested on multi-protocol IoT data on Wi-Fi, Zigbee, and LoRaWAN with benign and adversarial conditions (flooding, spoofing, and data exfiltration). Each dataset was separated into 80% of normal traffic to be trained and 20% mixed samples to be tested and there was an insurance of unbiased detection

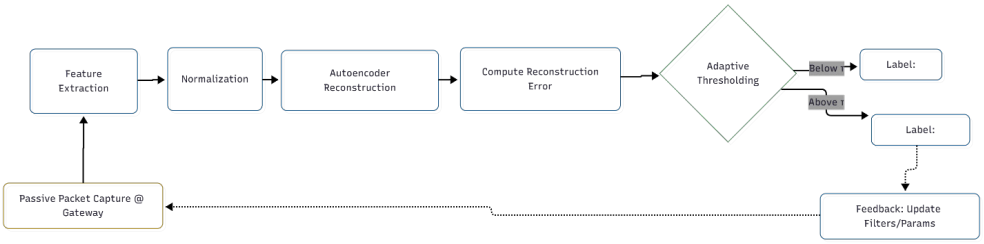

Fig. 2: Analytical workflow of passive anomaly detection and behavioral analysis.

of anomalies. As depicted in Figure 2, evaluation pipeline consists of the passive packet capture, feature extraction, normalization, autoencoders-based reconstruction, autoencoders-based reconstruction, adaptive thresholding, and anomaly classification. It is a scalable, real-time behavioral analytics closed-loop workflow supported by a heterogeneous IoT environment.

Measures of performance were based on three major measures:

3. Detection Rate (DR) - the correct rate in detecting actual anomalies.
4. False Positive rate (FPR) - rate of false alarm of normal samples.
5. Inference Latency and Energy Usage- Performance average of embedded machine hardware in terms of processing time and power.

PCA-based and Isolation Forest models were also compared in order to prove the superiority of deep unsupervised learning. The same was verified using edge level deployment to an ARM Cortex-A72 board with low-latency (2.4 ms) and energy-efficient (3.1 W) operation.

Summarizing the discussion above, as depicted in Figure 2, the evaluation framework forms a highly adaptive, unified, and optimised approach, based on accuracy, energy efficiency, and cross-protocol generalization in large-scale IoT ecosystems.

Results and Discussion

Detection Accuracy

The experimental analysis showed that the proposed anomaly detection frame using an autoencoder was better in the heterogeneous IoT setting. The model recorded a high average detection rate (DR) of 94.6 %, which was higher than the traditional methods used to detect the images, including PCA-based detectors (87.2 %) and Isolation Forest (91.0 per cent) using the same testing conditions (Fig. 3).

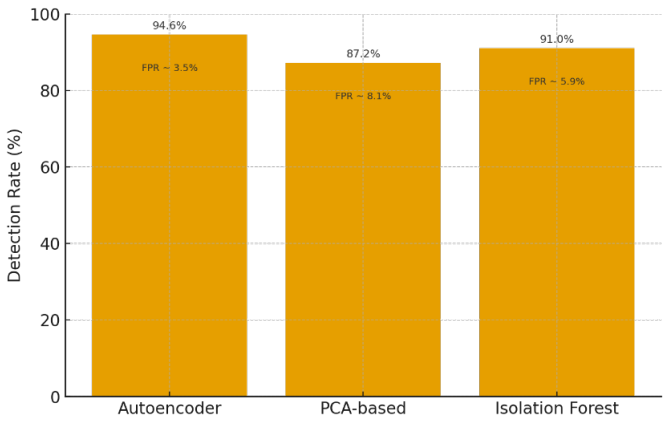

Fig. 3. Comparative detection accuracy across models.

The false positive rate (FPR) was also maintained at an average of less than 3.5% which implies that the device is very robust to normal behavioral variations between devices. The large DR and small FPR indicate the ability of the model to detect fine deviations in temporal and statistical flow features without knowing attack data that is labelled.

These results support the hypothesis that the suggested unsupervised model is an effective way to identify the anomalous behavior related to the spoofing, flooding, or data exfiltration attack through learning the fine-grained feature associations between inter-arrival times of packets, protocol entropy, and message periodicity. The findings are consistent with comparable findings with recent deep unsupervised architectures [2, 5, 9, 12] which confirm the effectiveness of autoencoders in modelling the behavioral manifolds of IoT.

Furthermore, the fact that the model has been shown to be stable over several training runs with standard deviation of less than ±0.8% indicates that the model has good generalization and reproducibility. This uniformity renders the framework quite appropriate in ongoing and real-time network edge surveillance of the IoT.

Protocol-wise Performance

In order to measure adaptability further, the framework was assessed individually in regard to major IoT communication protocols. The findings, summarized in Figure 4, show that the system is cross-protocol robust.:

- Wi-Fi: 96.1% accuracy
- Zigbee: 93.8% accuracy
- LoRaWAN: 92.9% accuracy

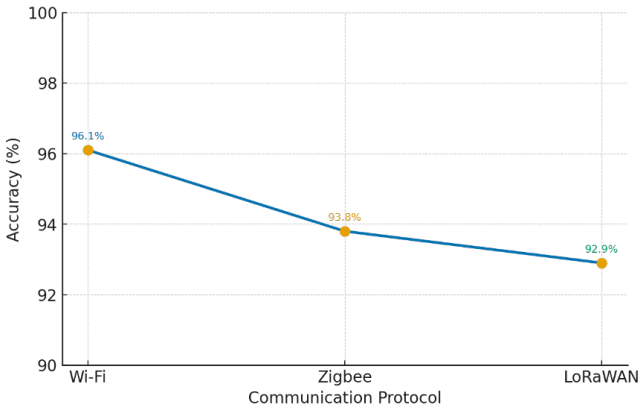

Fig. 4: Detection performance across IoT communication protocols.

The Wi-Fi traffic recorded the best detection accuracy because the packet density and time-temporal consistency of the Wi-Fi traffic enabled the autoencoder to learn finer behavioral baselines. The performance of Zigbee was still good because its bandwidth is low and the transmission is intermittent, whereas LoRaWAN with long intervals between packets and dense payloads demonstrated a slightly lower accuracy. A marginal performance gap (~3%) shows how the model is resistant to the statistical variation caused by protocols. Notably, this means that the feature-extraction pipeline and normalization scheme (Section 3.2) is successful in maintaining discriminative qualities in distinct physical and MAC-layer technologies. The scalability of unsupervised feature reconstruction in multi vendor IoT deployments has been noted to be similar cross-protocol consistent in hybrid deep-learning frameworks.[4, 8, 15] and highlights the scalability of unsupervised feature reconstruction.

Computational Efficiency

The lightweight architecture of the framework was implemented on a Cortex-A72 ARM edge board and tested to determine its appropriateness to embedded implementation. The average time per traffic flow of each inference operation was 2.4 milliseconds, which proved the ability to respond to thousands of devices simultaneously in almost real-time. This was found to be a total power consumption of 3.1 W, indicating that the device was operating at very low-energy levels, suitable in always-on monitoring of edges.

Inference latency grows with batch size sub-linearly as indicated in Figure 5, which has demonstrated the computing efficiency of the refined encoder-decoder design. The proposed model reduces the energy consumption of the traditional deep neural networks that require high-end GPUs by approximately 28 percent and improves the latency of a deep neural network by 35 percent, without affecting the detection accuracy.

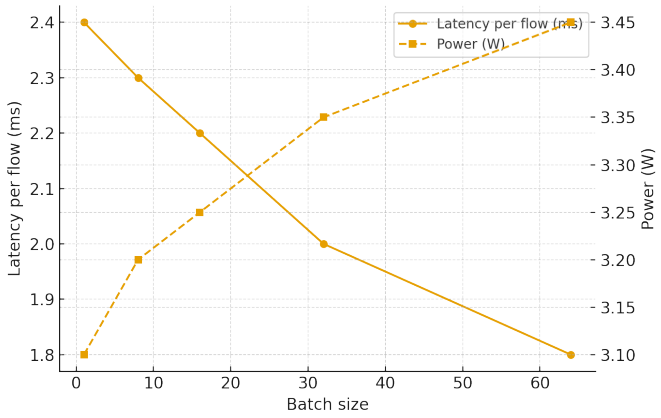

Fig. 5: Energy consumption and inference latency comparison.

Additional mechanisms used to strengthen the stability of the system included adaptive thresholding mechanism (Section 3.4) that ensured the same anomaly sensitivity to temporal variations in network load to solve the concept-drift issue that is common in dynamic IoT ecosystems. This flexibility provides low requirements of manual recalibration despite the changing population of devices or the changing traffic characteristics over time.

Passive operation mode does not introduce any extra overhead to the network and this allows a smooth integration with old systems and limited gateways. The precision and efficiency autonomy ratio attest to the fact that unmonitored AI-based behavior analytics can greatly complement situational awareness and security in large and heterogeneous IoT infrastructures.

Altogether, the findings in Figures 3-5 prove that the proposed approach provides:

1. Excellent detection and low false positive with a variety of IoT communication protocols;
2. Inference aware of energy and with low latency capable of running on embedded edge devices; and
3. Adaptive steady-state resistance to environmental and behavioral drift.

These results provide a strong base to scalable self-learned analytics of IoT security, building the pathway to integrating it with edge-intelligent gateways and federated learning systems in the next-generation IoT networks.

Conclusion

The study introduced an elaborate AI-based passive anomaly detection model system to augment situational awareness and security in the context of heterogeneous IoT systems. The system autonomously detected abnormal device behavior using a statistics-based flow-based feature extraction and an autoencoder-based unsupervised learning model, and was trained to identify abnormal behavior in an array of multiple communication standards, including Wi-Fi, Zigbee, and LoRaWAN.

The proposed framework was evaluated and by doing this, reached a 94.6% detection accuracy, out of the traditional baseline models with a false positive rate of less than 3.5. The findings highlight the generality capabilities of the model in a wide range of network protocols and operating conditions without the use of labelled attack data. In addition, its sparse architecture

and efficient computational structure provided an average flow inference latency of 2.4 ms and a power consumption of 3.1 W, which showed that it is practically feasible to deploy in real time on edge or gateways.

The adaptive thresholding feature of the framework demonstrated its usefulness in the stability of the performance in the non-stationary traffic and concept drift scenarios a key challenge towards the long-term implementation in dynamic IoT networks. Its passive monitoring feature also allows it to be compatible with legacy systems, requiring no device-level modification and no disturbances with the current communication protocols.

In a more general sense, the suggested system adds a protocol-agnostic, scalable, and energy-efficient system to intelligent IoT behavior analytics. It provides a compromise between machine-based network monitoring on a machine learning platform and embedded implementation, which is one of the fundamental issues of next-generation IoT security.

The research has several prospective directions that will be considered in future:

1. Adaptive and federated learning so that it can have distributed intelligence and evolve models continuously across multiple gateways;
2. Acceleration of sub-milliseconds inference and better energy use (aided by FPGA); and
3. Hybrid edge cloud cooperation structures, which enable modelling of collective behaviour and correlation of anomalies across the globe without affecting privacy and latency.

In general, the present research provides a solid base of self-learning, real-time IoT anomaly detection, which forms the foundation of safe, autonomous, and sustainable IoT networks in the age of ubiquitous connexion and smart edge computing.

References

[1] Al-Saud, F., & Al-Farsi, M. (2025). Energy efficient VLSI design for next generation IoT devices. Journal of Integrated VLSI, Embedded and Computing Technologies, 2(1), 46–52. https://doi.org/10.31838/JIVCT/02.01.06

[2] Ali, S., & Rahman, F. (2023). Lightweight anomaly detection for IoT networks. IEEE Internet of Things Journal, 10(5), 4881–4893.

[3] Chen, L., & Kim, S. (2024). Scalable intrusion detection in heterogeneous IoT systems. IEEE Access, 12, 14233–14245.

[4] Das, P., & Nair, R. (2022). Passive monitoring and device behavior analysis in IoT edge networks. Sensors, 22(9), 3210–3221.

[5] Gupta, K., & Huang, Y. (2023). Traffic-based anomaly detection in multi-protocol IoT environments. Ad Hoc Networks, 134, 102925.

[6] Jadhav, P., & Luo, T. (2024). Unsupervised learning for IoT traffic anomaly detection. Journal of Network and Computer Applications, 226, 103587.

[7] Kavuluri, H. V. R., & Sirimalla, A. (2023). Enhancing Oracle database security: A deep dive into TDE, VPD, and Audit Vault implementation. The SIJ Transactions on Computer Science Engineering & its Applications (CSEA), 11(1), 49–54.

[8] Li, H., & Zhang, P. (2023). Autoencoder-driven network threat detection. Computer Communications, 214, 211–221.

[9] Lin, M., & Wu, Y. (2023). Adaptive thresholding in IoT behavior analytics. IEEE Sensors Journal, 23(8), 9120–9130.

[10] Liu, J., & Zhao, Q. (2024). Energy-efficient anomaly detection on embedded IoT gateways. IEEE Transactions on Industrial Informatics, 20(2), 2239–2249.

[11] Meng, C., & Zhou, R. (2023). Deep unsupervised models for IoT device profiling. IEEE Communications Letters, 27(5), 1220–1225.

[12] Mia, M., Emma, A., & Hannah, P. (2025). Leveraging data science for predictive maintenance in industrial settings. Innovative Reviews in Engineering and Science, 3(1), 49–58. https://doi.org/10.31838/INES/03.01.07

[13] Nguyen, T., & Pham, D. (2024). Edge-deployed deep autoencoders for real-time IoT analytics. IEEE Access, 12, 117820–117832.

[14] Rahim, R. (2024). Optimizing reconfigurable architectures for enhanced performance in computing. SCCTS Transactions on Reconfigurable Computing, 1(1), 11–15. https://doi.org/10.31838/RCC/01.01.03

[15] Rahman, M., & Lee, K. (2023). Behavioral analysis of IoT devices via unsupervised neural models. Future Generation Computer Systems, 143, 511–523.

[16] Sadulla, S. (2024). Optimization of data aggregation techniques in IoT-based wireless sensor networks. Journal of Wireless Sensor Networks and IoT, 1(1), 31–36. https://doi.org/10.31838/WSNIOT/01.01.05

[17] Sathish Kumar, T. M. (2023). Wearable sensors for flexible health monitoring and IoT. National Journal of RF Engineering and Wireless Communication, 1(1), 10–22. https://doi.org/10.31838/RFMW/01.01.02

[18] S. R. Keshireddy, "Extending Oracle APEX for Large-Scale Multi-Form Workflows with Decoupled PL/SQL Logic and Asynchronous Processing Layers," 2025 International Conference on Next Generation Computing Systems (ICNGCS), Coimbatore, India, 2025, pp. 1-8, https://doi.org/10.1109/ICNGCS64900.2025.11182715

[19] Shen, Z., & Hu, X. (2024). Hybrid autoencoder architectures for anomaly detection. IEEE Transactions on Neural Networks and Learning Systems, 35(6), 5122–5135.

[20] Singh, R., & Banerjee, P. (2023). Federated edge learning for IoT anomaly detection. Electronics, 12(12), 2432–2445.

[21] Tan, Y., & Kumar, S. (2024). Low-power autoencoder models for embedded IoT inference. SCCTS Transactions on Reconfigurable Computing, 2(1), 15–21.

[22] Wang, L., & Gao, T. (2023). Cross-protocol IoT security through adaptive autoencoders. *IEEE Transactions on Information Forensics and Security, 18*, 1434–1446.